Rev. 23.0 Prime Networks Release Notes

RLN10252-1LA

June 1990

These Release Notes, RLN10252-1LA, contain seven chapters and one appendix updating the information in the following books. This information concerns file system name spaces and the Name Server process, network event messages, and miscellaneous PRIMENET information.

Manual

PRIMENET Planning and Configuration Guide
DOC7532-4LA, UPD7532-41A

Programmer's Guide to Prime Networks
DOC10113-1LA, UPD10113-11A

Operator's Guide to Prime Networks
DOC10114-1LA, UPD10114-11A

User's Guide to Prime Network Services
DOC10115-1LA, UPD10115-11A

RLN10252-1LA

Rev. 23.0 Prime Networks Release Notes

George W. Gove
Emily Stone

This manual documents the software operation of the PRIMOS operating system on 50 Series computers and their supporting systems and utilities as implemented at Master Disk Revision Level 23.0 (Rev. 23.0).

Prime Computer, Inc., Prime Park, Natick, MA 01760

The information in this document is subject to change without notice and should not be construed as a commitment by Prime Computer, Inc. Prime Computer, Inc., assumes no responsibility for any errors that may appear in this document.

The software described in this document is furnished under a license and may be used or copied only in accordance with the terms of such license.

Printing History

First Edition (RLN10252-1LA) June 1990 for Revision 23.0

Credits

Editorial: Mary Skousgaard
Project Support: Bill Huber, Stan Luke, Ron McKenzie
Illustration: Elizabeth Wahle
Production: Judy Gordon

How to Order Technical Documents

To order copies of documents, or to obtain a catalog and price list:

United States Customers	*International*
Call Prime Telemarketing, toll free, at 1-800-343-2533, Monday through Thursday, 8:30 a.m. to 8:00 p.m. and Friday, 8:30 a.m. to 6:00 p.m. (EST).	Contact your local Prime subsidiary or distributor.

PRIME SERVICE℠

Prime provides the following toll-free number for customers in the United States needing service:

1-800-800-PRIME

For other locations, contact your Prime representative.

Surveys and Correspondence

Please comment on this manual using the Reader Response Form provided in the back of this book. Address any additional comments on this or other Prime documents to:

Technical Publications Department
Prime Computer, Inc.
500 Old Connecticut Path
Framingham, MA 01701

CONTENTS

ABOUT THIS BOOK

These release notes, RLN10252-1LA, contain seven chapters and one appendix updating the information in the books listed in the section Related Documentation. This information concerns file system name spaces and the Name Server process, network event messages, and miscellaneous PRIMENET information.

ORGANIZATION OF THESE RELEASE NOTES

These release notes consist of seven chapters and one appendix, as follows:

- Chapter 1 discusses file system name spaces, use of file system object pathnames, and mounting of partitions in the file system hierarchy.
- Chapter 2 describes the considerations you use in planning a file system name space for a collection of networked systems.
- Chapter 3 discusses the network considerations you use, including aspects of remote file access (RFA) and forced user validation (FUV), and briefly discusses two example networks.
- Chapter 4 discusses configuring the network, DSM, and DSM unsolicited message handling (UMH) in order to create a common file system name space. It also describes the command for starting the Name Server.
- Chapter 5 discusses operation of the Name Server, including monitoring the common file system name space.
- Chapter 6 describes the network event messages that are produced by PRIMENET and logged by the Distributed Systems Management (DSM) facility.
- Chapter 7 includes miscellaneous PRIMENET information related to START_NM and STOP_NM commands, the START_NET command, gateway configuration guidelines, ACL group restriction for RFA, and XLCONN and XLASGN subroutines.
- Appendix A describes Name Server messages reported by DSM.

RELATED DOCUMENTATION

These Release Notes update the information in the following books:

- PRIMENET Planning and Configuration Guide (DOC7532-4LA and UPD7532-41A)
- Programmer's Guide to Prime Networks (DOC10113-1LA and UPD10113-11A)
- Operator's Guide to Prime Networks (DOC10114-1LA and UPD10114-11A)
- User's Guide to Prime Network Services (DOC10115-1LA and UPD10115-11A)

For other information related to common file system name spaces, changes in pathname syntax, accessing file system objects, subroutines related to the file system, the Name Server process, and commands related to the new file structure and the Name Server process, see these books:

- System Administrator's Guide, Volume I: System Configuration (DOC10131-3LA)
- Operator's Guide to System Commands (DOC9304-5LA)
- PRIMOS User's Release Document (DOC10316-1PA)
- Advanced Programmer's Guide II: File System (DOC10056-3LA)
- Subroutines Reference II: File System (DOC10081-2LA)
- DSM User's Guide (DOC10061-3LA)

To obtain a complete list of Prime technical documentation online, type the command HELP DOCUMENTS. A hardcopy list is available in the *Guide to Prime User Documents.* This guide is issued twice a year. Lists of additional updated material are published quarterly in the Customer Service Newsletter.

PRIME DOCUMENTATION CONVENTIONS

The following conventions are used throughout this document. The examples in the table illustrate the uses of these conventions.

UPPERCASE

In command formats, words in uppercase bold indicate the names of commands, options, statements, and keywords. Enter them in either uppercase or lowercase.

Example

ADD_PORTAL

italic

In command formats, words in lowercase bold italic indicate variables for which you must substitute a suitable value. In text and in messages, variables are in non-bold lowercase italic.

Example

ADDISK *pdev*

pdev is the physical device number.

Brackets

Brackets enclose a list of one or more optional items. Choose none, one, or several of these items.

Example

MAKE [-DISK / -PART]

Underscore

In examples, user input is underscored but system prompts and output are not.

Example

```
OK, ADDISK 4260
Starting up revision 23.0 partition "CMDDSK "
OK,
```

Hyphen

Wherever a hyphen appears as the first character of an option, it is a required part of that option.

Example

ADD_PORTAL -HELP

1

FILE SYSTEM NAME SPACES

THE PRIMOS FILE SYSTEM

At Rev. 23.0, the logical structure of the PRIMOS file system has changed. The new file system name space is tree-structured with a single root directory designated by < (the less-than symbol or left angle-bracket). The root directory is special in that it contains only other directories, or root entries. Directories other than the root directory can contain other file system objects including other directories and files.

A **name space** is the collection of the file system object names on systems in a DSM configuration group. The Name Server process manages this collection.

Note

For more detailed discussions of the concepts presented in this chapter, refer to related documentation listed in About This Book.

Physical Structure

The file system is stored on physical disks associated with either the local system or with remote systems. These disks are subdivided into logical partitions and each partition has a tree-structured file system beginning with a top-level directory known as the MFD.

At Rev. 23.0, each partition must be grafted, or mounted, either into the root directory or over a directory of another local partition that has already been added. Partitions can be mounted at any directory level within the tree structure.

Directories in the root directory can be physically located either on the local system or on remote systems. If the Name Server process is running, the root directory contains an entry for each disk on all the systems in a particular DSM configuration group, unless the disk is mounted lower in the file system. The Name Server processes on the systems in that DSM configuration group create the root directory by replication of their local disk lists. In this way, the file systems of the members of the DSM configuration group becomes a **common file system name space.**

The list of partitions that comprise the common file system name space at Rev. 23.0 is stored in a table known as the **Global Mount Table (GMT).** The contents of the GMT include the names of all partitions and where the partitions are mounted within the file system (their mount-point pathnames). You use the LIST_MOUNTS command to see the contents of the GMT.

IDENTIFYING OBJECTS IN THE FILE SYSTEM

An object in the tree-structured file system name space is identified by its pathname. A pathname contains the sequence of directories that must be traversed in order to get to a particular file system object. The common file system name space is defined and bounded by the file systems of the members of a particular DSM configuration group. In the common file system name space, the root directory of all systems sharing the name space is identical.

Pathname Syntax and Semantics

At Rev. 23.0, there are new forms of pathname syntax. These new forms are

- < (the less-than symbol or left angle-bracket) for the root directory
- <ENTRY, representing a logical partition, for an entry, or directory, in the root

Pathname interpretation has been modified so that the root is the starting point in the file system name space. Previously, the starting point was the disk table on a system, or the list of added disks for a system. You can attach to the root directory and list its contents as with any other directory. The ACL on the root directory is LU and cannot be changed. Thus the root cannot be changed except by the Name Server.

The entries, or directories, in the root, which correspond to partitions, can be referenced either as <DISK1 or as <DISK1>MFD. The two forms are synonymous. < is no longer a delimiter in a pathname; < indicates the root directory.

PRIMOS subroutines have been changed to recognize the new forms of pathnames.

Pathname Categories

There are two categories of pathnames, fully-qualified and unqualified. **Fully-qualified** pathnames specify a complete, unambiguous path for locating an object in the file system tree structure. Fully-qualified names begin either with < (the root) or with *> (relative to the current attach point).

Unqualified pathnames specify a partial pathname and require PRIMOS to search through the ATTACH$ search rules to find an object.

Interpreting Unqualified Pathnames

Unqualified pathnames use the ATTACH$ search rules to determine which objects to reference. ATTACH$ search rules define a list of pathname prefixes, usually a list of partition names, that are prepended to the unqualified pathname to make a fully-qualified pathname.

An attempt is made to attach to the first prefix and then operate on the first component in the unqualified pathname. If the component does not exist in that directory, each prefix in the list is tried until one is found that contains the desired component.

Before Rev. 23.0, the prefixes had to be a disk name. At Rev. 23.0, the prefix can be the pathname of any directory.

-ADDED_DISKS is a special ATTACH$ search rule that allows searching of the entire list of disks known to the local system. At Rev. 23.0, the interpretation of -ADDED_DISKS depends on whether the Name Server is running:

- If the Name Server is *not* running, the local disk table determines the list of disks searched. Local disks are searched first and then remote disks, both in the order they are found in the table.
- If the Name Server *is* running, the GMT determines the list of disks searched. Local disks are searched first (in ldev order) and then remote disks are searched in an internally optimized order.

If the Name Server is running such that there is a common file system name space among systems, you may want to use explicit ATTACH$ search rules, particularly in these cases:

- If there are many disks in the common file system name space, searching the GMT may take longer than searching the disk table, since the GMT has many more entries.
- If the attach scan performance becomes an issue, add explicit search rules to the ATTACH$ search rules. For example, add the names of the disks searched most often to the top of the list.

Note

Disks mounted within other disks, that is, not in the root but lower in the file system (as discussed in a following section), are not searched as part of the -ADDED_DISKS search rule. To be searched, their pathnames must be entered explicitly in the ATTACH$ search rules.

Using Different Pathname Forms in Applications

Fully-qualified pathnames provide better performance than unqualified pathnames since no searching is necessary. Fully-qualified pathnames unambiguously identify objects. Thus, fully-qualified pathnames eliminate problems due to programs depending on the disk table order or on the ATTACH$ search rules.

Because of logical disk mounts at Rev. 23.0, in which partitions can be mounted anywhere in the file system, pathnames of objects do not have to change if applications expand to additional partitions.

With the Name Server running, distributed applications can use fully-qualified pathnames and be guaranteed that the same object is referenced regardless of which system in the DSM configuration group is running the program.

Using unqualified pathnames allows objects to be referenced with shorter names. In addition, because the name is not qualified by a partition name, the application can use alternate top-level directories if a disk or system becomes inaccessible, assuming use of an alternate directory is acceptable.

MOUNTING PARTITIONS INTO THE FILE SYSTEM HIERARCHY

At Rev. 23.0, the ADDISK command has a new option, -MOUNT_PATH, that specifies where a disk is to be added in the file system hierarchy. (See the *Operator's Guide to System Commands* for a complete discussion of ADDISK.)

If you do not specify the -MOUNT_PATH option, the partition is added to the root directory. Adding a partition to the root means a directory, or root entry, is created in the root directory. If the -MOUNT_PATH option is not used, the name of the directory is the same as the name of the partition. If a user does a relative attach from the root to the newly created directory (root entry), the user attaches to the MFD of the newly added partition.

If you use the -MOUNT_PATH option, you can specify the mount point as either a directory in the root or as any existing directory on an already added partition. Figure 1-1 shows three examples of adding a partition in various places in the file system hierarchy and these examples are explained following the figure.

```
1.  Command:  ADDISK 4062

    Result:   <PARTN

2.  Command:  ADDISK 4062 -MOUNT_PATH <NEW_DIRECTORY_IN_ROOT

    Result:   <NEW_DIRECTORY_IN_ROOT

3.  Command:  ADDISK 4062 -MOUNT_PATH <BIGDSK>TOPDIR>SUBDIR

    Result:   <BIGDSK>TOPDIR>SUBDIR
```

FIGURE 1-1. Adding Partitions With Various Mount Paths

If you do not use the -MOUNT_PATH option, the partition is added in the root with the name given it by MAKE as in example 1 of Figure 1-1. If you mount the partition in the root using the -MOUNT_PATH option, a directory is created in the root with the given name, for example, <NEW_DIRECTORY_IN_ROOT as in example 2. Since the mount point representing the partition is a directory name, it can be up to 32 characters long, for example, <A_REALLY_LONG_DIRECTORY_NAME.

Caution

Pre-Rev. 23.0 systems cannot access partitions added in the root with the -MOUNT_PATH option using a partition name longer than six characters. If you have a mixed revision environment, you should not add partitions in the root with names longer than six characters.

When a partition is added with the -MOUNT_PATH option, the name of the mount-point directory must not already exist in the root. The name of the partition, as given when it was created with MAKE, does not appear in pathnames which identify objects on the newly added partition; the mount-point directory name appears instead.

Mounting Partitions Lower in the File System Hierarchy

If you use the -MOUNT_PATH option to add a partition at a mount-point directory below the root directory, the mount-point directory must already exist and that directory can be at any level in the file system hierarchy, for example, <BIGDSK>TOPDIR>SUBDIR as in example 3 of Figure 1-1. User SYSTEM (User 1) must have a minimum of Use (U) access to the mount-point directory in order to add the partition at that point.

After you add the partition in this way, attaching to the specified mount-point directory (<BIGDSK>TOPDIR>SUBDIR) puts you in the MFD of the newly added partition.

The previous contents of the mount-point directory (SUBDIR) become inaccessible to users not already attached there or to the parent directory as long as the partition remains added in this manner. Therefore, if you want the contents of the mount-point directory available, you must transfer the contents to the partition added to the mount-point directory using a procedure such as that discussed in *System Administrator's Guide, Volume 1: System Configuration.*

The ACLs governing access to directories on the partition mounted in this way are determined first by the ACLs protecting the mount-point directory and then by the ACL on the MFD of the newly added partition.

Note

The backup products (DBR) do not cross mount points. Therefore partitions accessible through a mount point are not backed up when you back up the partition that contains the mount point. The contents of the mount-point directory is not backed up either, because the contents are inaccessible. Thus you should back up the mount-point directory before mounting a partition there, or you should use an empty directory as the mount-point directory.

An advantage of adding partitions lower in the file system and that are subordinate to other partitions is that it is easier to extend the storage capacity of systems whose applications are written to use fully-qualified pathnames. By replacing large directories with a new partition, the pathnames needed by the applications remain unchanged.

The ability to mount subordinate partitions also means that there is no limit to the size that a directory can grow. In addition, mounting subordinate partitions eliminates the requirement of having hundreds of entries at the top of the file system hierarchy (in the root directory) on networks containing hundreds of partitions.

2

PLANNING FILE SYSTEM NAME SPACES

PLANNING A COMMON FILE SYSTEM NAME SPACE FOR A COLLECTION OF NETWORKED SYSTEMS

PRIMENET allows users to transparently access file system objects without the user needing to know on which system the file system objects reside. At Rev. 23.0, the file system can be built as a collection of systems that share a common file system name space, that is, by using the Name Server. The only way to build a file system name space in a network before Rev. 23.0 was on a system-by-system basis.

If the file system name space is built on a per-system basis, as on pre-Rev. 23.0 systems,

- Each system has its own unique list of added disks that appear in the disk list and the list is different than the lists of other systems in the collection.
- The disk list must be created manually by using ADDISK commands for both local and remote disks.

Figure 2-1 illustrates name spaces built on a per-system basis as on pre-Rev. 23.0 systems. Each system has its own unique disk list that is manually constructed by using the -ON option of ADDISK.

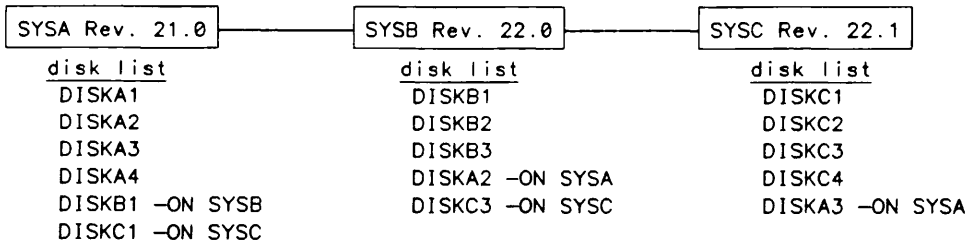

FIGURE 2-1. File System Name Space - pre-Rev. 23.0

If the file system name space is built as a collection of systems,

- Some collection of systems is defined as sharing a common file system name space.
- At Rev. 23.0, the collection of systems consists of those systems that are in the same DSM configuration group.
- Having a common file system name space means the root directory on all these Rev. 23.0 systems in the same DSM configuration group is identical.
- The mechanism that replicates the root directory so that it is identical on all these systems is a server process known as the **Name Server.**
- If the Name Server is running on each of the systems in the DSM configuration group, disks on remote systems in the DSM configuration group (except those disks that are mounted lower in the file system) are automatically added to the root directory. In this way, all systems that share in the common file system name space see all disks. Lower mounted disks appear in the GMT.
- If there are pre-Rev. 23.0 systems in the DSM configuration group, there is no Name Server process running on those systems and they are polled by the Rev. 23.0 Name Server processes to obtain their local disk lists for addition to the root directory.

Figure 2-2 illustrates a collection of systems sharing a common file system name space. The Name Server process runs on the Rev. 23.0 systems (SYSA and SYSC) and creates the root and the GMT for those systems by broadcasting the local disk lists to the other Rev. 23.0 systems and by polling the pre-Rev. 23.0 system (SYSB). In this way, as soon as the Name Server process starts on SYSA and SYSC, those systems have identical root directories and the GMT is common between them. The pre-Rev. 23.0 system, SYSB, must have remote disks manually added but its local disk list is obtained by the Rev. 23.0 Name Server processes on SYSA and SYSC polling SYSB and those disks become part of the common file system name space. If SYSB is later converted to Rev. 23.0 and the Name Server process is started on SYSB, it will have a root directory identical to those on SYSA and SYSC and will share the same GMT.

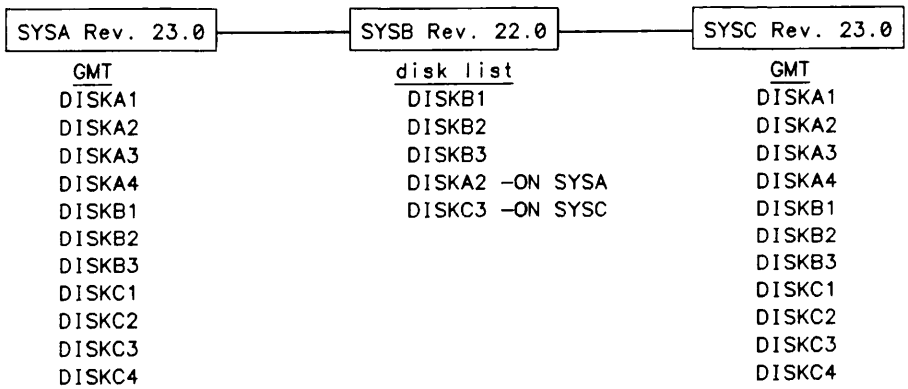

FIGURE 2-2. File System Name Space - Rev. 23.0

Deciding How to Build the File System Name Space

As discussed above, there are two alternative means for building the file system name space, namely, as a collection of systems or on a per-system basis. You should consider some of the following reasons for wanting a common file system name space among a collection of systems:

- You want to view a particular collection of systems in the network as a single computing resource because users or applications reference file system objects on multiple systems and the systems do not impose forced user validation (FUV) on each other.
- You want to eliminate the administrative overhead of having to maintain remote disk lists on each system. Disks that are added to one system are automatically added to all other systems that share the common file system name space; therefore, use of ADDISK commands for remote disks are not needed.

  The Name Server on each system also tracks the addition of new or moved disks (including dual-ported disks) and updates all the other systems in the common file system name space. Automatically updating the disk information makes it easier to move disks to alternate systems when there are hardware failures. In addition, the root directory is thus dynamically maintained with disks being added to and removed from the root as they are added to and removed from systems in the common file system name space.
- You want to develop or run distributed applications. Fully-qualified pathnames uniquely identify file system objects regardless of which system the reference is made from because the Name Server guarantees that the root directories on all systems in the DSM configuration group are identical. Unique names mean programs can run on any system and still get the same result because the same objects are always referenced.

  Being able to run programs on any system means it is easier to distribute computations in a networked environment and it is easier to move programs to backup or alternate systems when there are hardware failures.
- You want to have up to 1280 local and remote disk partitions added to each system versus the 238 available without the common file system name space.

Some of the reasons for wanting to configure the file system name space on an individual system basis include

- There is forced user validation (FUV) between systems. It is less desirable to have a common file system name space if the network environment is not open.
- The systems chosen for joint administration by being put in the same DSM configuration group do not want to share a common file system name space.
- There are different disks on different systems that have the same name and these names cannot change, precluding addition of these disks to the root with the existing name of the disk.

You must carefully plan the collection of systems to share a common file system name space because of the assumptions the Name Server makes about PRIMOS revisions. You must decide how to partition an existing network into one or more common file system name spaces considering the following:

- Each system can be in only one name space, that is, name spaces cannot overlap.
- The systems in the common file system name space can be at any supported PRIMOS revision but only Rev. 23.0 systems can run the Name Server.

Starting the Name Server

You must also plan when to first start the Name Server on each system in the DSM configuration group. The Name Server polls pre-Rev. 23.0 systems for their disk lists but does not poll Rev. 23.0 systems because each Rev. 23.0 system's Name Server assumes a Name Server is running on every other Rev. 23.0 system in the DSM configuration group.

Thus, if a Name Server is started on one system in the DSM configuration group, a Name Server should be started on each system in the same DSM configuration group that is running PRIMOS Rev. 23.0. In addition, once one system in a DSM configuration group is running the Name Server, any other systems in the same DSM configuration group that are later converted to Rev. 23.0 should also run their Name Server.

Operation of the Name Server in Mixed Revision Environments

A DSM configuration group may contain systems that are running Rev. 23.0 and pre-Rev. 23.0 versions of PRIMOS. The Name Server processes running on the Rev. 23.0 systems poll the pre-Rev. 23.0 systems to obtain information about the disk lists on those pre-Rev. 23.0 systems, that is, the list of local disks on those systems. The Name Servers poll at specific time intervals (the default is 10 minutes) to detect any changes in the local disk lists of those systems due to ADDISK or SHUTDN commands. In this way, the Name Servers update the root directories and the GMT of the Rev. 23.0 systems in the DSM configuration group.

The pre-Rev. 23.0 systems in the DSM configuration group have a limited view of the common file system name space because remote disks must be added manually to those systems. When those systems are later converted to Rev. 23.0, it is only necessary to start the Name Server process on them because they are already a part of the DSM configuration group and can thus share in the group's common file system name space.

Adding Remote Disks in Relation to the Name Server

In general, it is not necessary to add remote disks to a system by using the ADDISK command if the Name Server is running. When the Name Server is running, remote disks on systems in the same DSM configuration group are automatically added to the root directory, allowing them to be referenced. Remote disks in other name spaces can be accessed through portals, which are discussed below.

The Name Server does not, however, add remote disks to the local disk table. The only disks in the local disk table (as shown by a STATUS DISKS or LIST_DISKS display) are the local disks as shown in this example.

```
OK, STATUS DISKS

                                        Mirror
                                        Primary Secondary  State
Disk    Ldev  Pdev  System  Robust    ----------------------------------
REGRP1   0   6060                        6060      6162  Active
REGRP2   1   5222
REGRP4   2   5523
REGRP5   3   5621
REGRP9   4  52420

Note: The Name Server is started on your system. Therefore, the STATUS
      DISKS command may not display all the disks to which you have
      access. To see the complete list, use the LIST_MOUNTS command.
```

Therefore, remote disks do not have logical device numbers (ldevs).

If there are programs that reference file system objects on remote disks by using ldev pathname syntax (for example, <5>DATABASE), there must be a way for PRIMOS to convert ldevs to partition names for compatibility. To accomplish this, add the remote disk with the ADDISK command after the Name Server is running so that the disk is added to the local disk table and thus is given an ldev. PRIMOS can then map the ldev to a partition name.

The remote disk must, however, already be added to the common file system name space root directory by the Name Server to enable access to it. Just being in the local disk table does not mean it can be referenced. When you add a disk in this manner, PRIMOS warns you that this is the case.

```
OK, ADDISK REGRP1 -ON SYSA
Warning: Remote addisks after starting Name Server are only used to
         provide ldev number to disk name mapping. The list of disks
         in the file system name space is given by the global mount
         table.
OK,
```

Note

The Name Server will not start if you add remote disks to the common file system name space by using the -ON option of ADDISK before you start the Name Server.

Use of Portals and Private Partitions in the Rev. 23.0 File System Name Space

The following means, in addition to use of the Name Server, may be used to construct the file system name space:

- Portals
- Private partitions

These means of constructing a file system name space are available at Rev. 23.0 whether or not you are running the Name Server process.

Portals: Portals are a new feature at Rev. 23.0. Portals allow users to be able to reference file system objects outside of their current file system name space. As shown in Figure 2-3, a portal is created from SYSA in Name Space X to SYS1 in Name Space Y, potentially allowing users on any system in Name Space X to access file system objects on any system in Name Space Y.

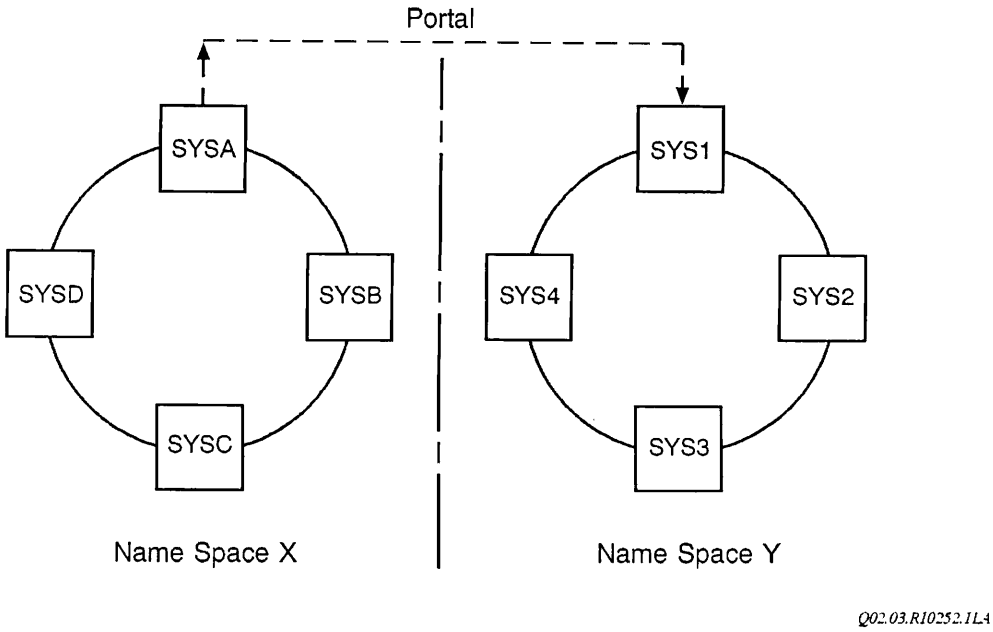

FIGURE 2-3. Accessing Another Name Space via a Portal

A **portal** is a directory on the user's system that has been transformed so that references to it are redirected to a directory on a remote system, the **target directory.** The target directory of a portal depends on the portal type:

- **Root-directed portals** redirect references to the root directory of another Rev. 23.0 file system name space.
- **Disk-directed portals** redirect references to the MFD of the specified disk on a remote system in another file system name space.

Root-directed portals are more powerful than disk-directed portals because root-directed portals potentially provide access to tens of systems and hundreds of disks. Disk-directed portals exist for compatibility with pre-Rev. 23.0 systems that do not have a root directory and, therefore, the portal must be to a disk.

Portals are a one-way means of access, that is, a portal in your file system name space directed to another file system name space allows you access to the other name space but does not allow access from that name space back to your name space. In addition, users in

your name space on pre-Rev. 23.0 systems can not use portals. (If they do, they receive error code 354, which is related to multiple hops between systems.)

Portals are created by using the ADD_PORTAL command and are removed by using the REMOVE_PORTAL command. (For a discussion of these commands and their options, see the *Operator's Guide to System Commands.*)

Portals increase the number of file system objects visible to users in the same manner as the ADDISK command does but portals are much more powerful in these ways:

- Adding a root-directed portal effectively grafts many systems and many disk partitions into the current file system name space.
- All Rev. 23.0 systems in the common file system name space can access a portal and, thus, the root of the name space that the portal is directed to.
- The local system can set the ACLs to the portal and, thus, set the access to the remote disk.

If there is forced user validation (FUV) between the local system and the remote system that is the target of the portal, local users must have logins on the target system and the local users must add remote user IDs with the ADD_REMOTE_ID (ARID) command before user access is allowed. Chapter 3, Network Considerations, discusses FUV and remote IDs. (See the *PRIMOS Commands Reference Guide* for details on ARID.)

If the volume of traffic over a portal is significant, you may want to include the target system in the current file system name space.

Portals are shown in the LIST_MOUNTS display as follows:

```
[LIST_MOUNTS Rev. 23.0 Copyright (c) 1990, Prime Computer, Inc.]
Mount   System Disk   Mount
 type   name   name   pathname
-----------------------------------------------
portal  ENGRGA        <BABEL>TCG>LAB ==> ENGRGB<DRWNG>
portal  ENGRGZ        <MECHAN>SYSCAD ==> CADCAM
portal  ENGRG1        <ELECTR>POWER>BACKUP ==> ENGRG2<BACKUP>
```

In the LIST_MOUNTS display,

- The `Mount type` is either `disk` for a disk mounted in the common file system name space or `portal` for either a root-directed or a disk-directed portal to another name space.
- `System name` is the system on which the portal was created.
- `Disk name` is blank when the mount type is a portal, otherwise it is the name of the disk at the mount point.
- `Mount pathname` is the fully-qualified pathname of the mount point directory on the system where the portal is created.
- The name immediately to the right of the arrow (==>) is the name of the system to which the portal is directed.

- The name that is in angle brackets (< >) to the right of the system name to which the portal is directed is the name of the disk to which all references to a disk-directed portal go. This field is blank for a root-directed portal.

For example, in the first portal listed, the Mount type is a portal from System name ENGRGA to the Disk name DRWNG with a Mount path on ENGRGA of <BABEL>TCG>LAB (the mount point is the directory LAB) to the system named ENGRGB and this is a disk-directed portal. Attaching to the pathname <BABEL>TCG>LAB on ENGRGA attaches the user process to the disk DRWNG on ENGRGB. Thus, this portal potentially provides access to the file system objects on one disk.

In the second portal listed, the Mount type is a portal, from the System name ENGRGZ with a Mount path on ENGRGZ of <MECHAN>SYSCAD to the system named CADCAM. Because there is no disk name after the name of the system to which the portal is directed, this is a root-directed portal. Attaching to the pathname <MECHAN>SYSCAD on ENGRGZ attaches the user process to the root directory of the file system name space that the system CADCAM is in. Thus, this portal potentially provides access to all the file system objects in another common file system name space.

The ACLs governing access to portals are determined first by the ACLs protecting the parent of the mount-point directory and then by the ACLs on the target of the portal.

Note

The backup products (DBR) do not cross mount points. Therefore, partitions accessible through a portal are not backed up when you back up the partition that contains the portal. The contents of the mount-point directory is not backed up either, because it is inaccessible. Thus you should back up the mount-point directory before mounting a partition there, or you should use an empty directory as the mount-point directory.

Private Partitions: Private partitions are another new feature at Rev. 23.0. Private partitions are disk partitions added with the -PRIVATE option of ADDISK. Private partitions cannot be accessed by users on remote systems with remote file access (RFA) using NPX slaves. (RFA is implemented internally with the network process exchange (NPX) facility.) Private partitions make the common file system name space, that is, use of the Name Server, useful in more situations. For example,

- All disks are visible from all systems with a common file system name space. If a site wishes to configure their systems in a common file system name space but wishes to keep a few disks that are to be referenced locally only, these disks can be added as private partitions.

The visibility of private partitions depends on where they are mounted in the file system:

- The mount-point directories of private partitions that are mounted in the root directory are visible to all users in the common file system name space through the use of the LD command on the root directory because of the LU ACL on the root, but cannot be referenced by remote file access (RFA).
- The visibility of private partitions mounted somewhere in the file system other than in the root directory depends on the ACL on the directory containing the mount point.

- Private partitions on remote systems do not appear in the display of the LIST_MOUNTS command unless the user is either at the supervisor terminal or is logged in as the System Administrator.

Although private partitions are not accessible with RFA, they are listed in the root directory if they are added in the root. Thus, partitions with the same name cannot be added to different systems as private partitions because a name conflict results. The purpose of private partitions is simply to prevent RFA access.

3

NETWORK CONSIDERATIONS

In planning how to divide your network into file system name spaces, you should consider where remote file access (RFA) and forced user validation (FUV) are configured in the network. Common file system name spaces, RFA, and FUV are related in important ways. This section provides guidelines for coordinating common file system name space configuration with the assignment of RFA and FUV in your network. In following the guidelines provided here, you may decide to define your name spaces based on the patterns of RFA and FUV that already exist in your network; or you may decide to change the assignment of RFA or FUV between some systems.

BRIEF REVIEW OF RFA AND FUV

RFA is a PRIMENET service that allows users to access file system objects on remote systems. RFA is implemented internally with the network process exchange (NPX) facility. RFA is configured between systems by means of CONFIG_NET, PRIMENET's network configuration utility. CONFIG_NET ensures that RFA is always configured symmetrically. That is, if you grant SYS1 RFA access to SYS2, CONFIG_NET automatically grants SYS2 RFA access to SYS1.

FUV is a security feature that requires users to add remote IDs before accessing remote file system objects. If SYS1 imposes FUV on SYS2, then a SYS2 user who wants to access a file on SYS1 must first use the ADD_REMOTE_ID (ARID) command to define a remote ID on SYS1. FUV is also configured by means of CONFIG_NET, and is usually configured symmetrically between two systems. The use of FUV between two systems is recommended if

- Security between the systems is a strong concern.
- The systems do not coordinate user IDs, so that IDs are not unique across the two systems. (The same ID might exist on both systems, representing a different user on each system.)

FUV is not necessary if

- Security between the systems is adequately handled by ACLs.
- User IDs uniquely identify users across the two systems. (If the same user ID exists on both systems, it represents the same user on both systems.)

RFA and FUV are configured independently of one another. If you have configured RFA between two systems, you can use the above recommendations to decide whether to configure FUV between them. The *PRIMENET Planning and Configuration Guide* discusses configuration of RFA and FUV in more detail.

COORDINATING RFA AND FUV WITH COMMON FILE SYSTEM NAME SPACES

The systems in a common file system name space generally have these characteristics:

- They are in the same DSM configuration group (by definition).
- They have general access to file system objects on all systems in the group, that is, RFA is enabled throughout.
- They may be managed by the same administrator.
- There is mutual trust among them and security is adequately handled by ACLs.
- User IDs are coordinated to be unique across all systems.

Systems in a common file system name space are usually managed by one administrator. FUV should not be necessary within a common file system name space because the System Administrator can coordinate user IDs. Although FUV may be configured between Rev. 23.0 systems with Name Servers, it is recommended that you do not configure FUV between systems. ACLs can then provide sufficient protection for files, so that RFA can be configured with confidence among all systems in the common file system name space.

On the other hand, systems in different DSM configuration groups do not share the same file system name space. To view or access files across the file system name space boundary, users must use portals. The systems may be managed by different System Administrators who do not coordinate user IDs. Security may be a strong concern between the systems. RFA can be configured, but FUV should be configured to prevent user ID impersonation if there is no coordination of user IDs.

Follow these guidelines for proper implementation of the common file system name space.

- It should not be necessary to configure FUV between systems in the same common file system name space. If you do configure FUV, configure it only between systems in different common file system name spaces.
- Configure RFA between systems in the same file system name space so that all users have access to the objects in the common file system name space.

- Within a common file system name space, all systems must be directly connected with no intervening (gateway) systems, because RFA is not supported over a gateway link.
- You may choose whether to configure RFA between systems in different common file system name spaces. If you do configure RFA between name spaces, portals must exist to enable users to access files across the name space boundary.

A SAMPLE NETWORK

Consider the network shown in Figure 3-1.

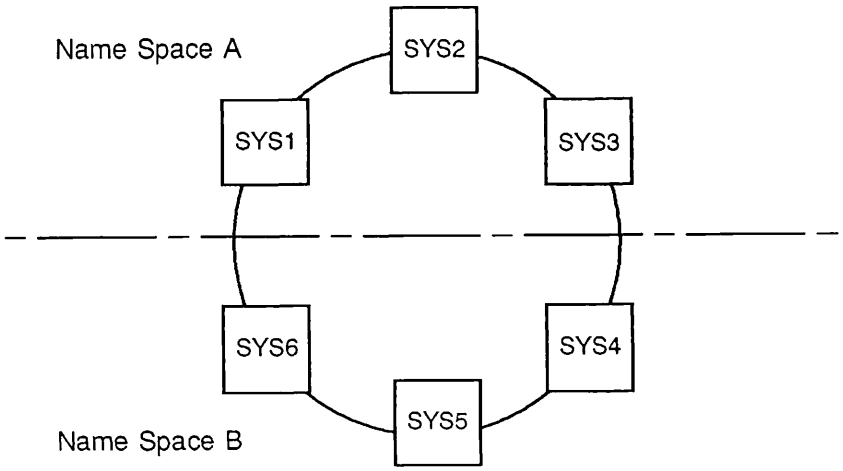

FIGURE 3-1. Two Name Spaces in a Networked Environment

The Rev. 23.0 systems in Figure 3-1 are in the same ring network and RFA is configured among all systems. Systems SYS1, SYS2, and SYS3 are in one DSM configuration group and systems SYS4, SYS5, and SYS6 are in another DSM configuration group. These systems comprise two common file system name spaces, Name Space A and Name Space B. All users in Name Space A can see all file system objects on the systems in Name Space A and all users in Name Space B can see all file system objects in Name Space B.

If an operator in Name Space A creates a root-directed portal from SYS3 to SYS4, users in Name Space A (systems SYS1, SYS2, SYS3) can list all entries in the root of Name Space B and can access files in Name Space B, depending on ACLs. Users in Name Space B, however, do not have access to Name Space A unless a portal is created from B to A.

Now consider the networks in Figure 3-2, which consist of a ring network with systems SYSA, SYSB, SYSC and SYSD and a LAN300 with systems SYSE, SYSF, SYSG, and SYSH. These two networks are connected by a synchronous line.

The systems within each of the networks have RFA enabled among them but there is no RFA enabled between the two networks. The synchronous line may be for FTS. There may be security considerations between the two networks or they may simply be physically remote from one another.

These two networks follow the rules for establishing common file system name spaces such as Name Space A and Name Space B as shown in the figure. The two separate name spaces could each be subdivided into other name spaces.

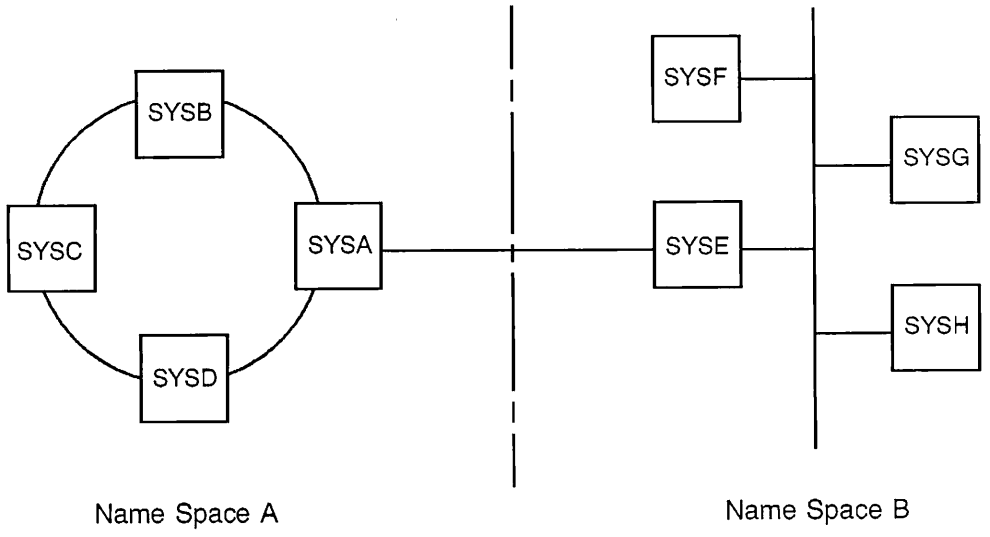

FIGURE 3-2. Name Spaces Associated with Different Networks

4

CONFIGURATION PROCEDURES

CREATING A COMMON FILE SYSTEM NAME SPACE AMONG SYSTEMS IN A NETWORK

To create and use a common file system name space, you need to

- Configure the network so that RFA is enabled between systems and FUV is not configured between systems in the common file system name space.
- Use DSM to define which systems share a common file system name space.
- Put the START_NAMESERVER command in the PRIMOS.COMI startup file.

Note

If you have upgraded from a pre-Rev. 23.0 system, you may also need to remove from your PRIMOS.COMI file any ADDISK commands that add remote disks.

Configuring the Network

Configure the network using CONFIG_NET based on the considerations you made concerning RFA and FUV and the file system name space boundaries. For details on using CONFIG_NET, see the *PRIMENET Planning and Configuration Guide.*

Using DSM to Define the Common File System Name Space

The list of systems in the DSM configuration group determines the boundaries of the common file system name space, that is, the list of systems that share disks. Before the Name Server is started, the DSM configuration groups on all these systems must be identical and DSM must be running on each system. This is the typical situation if DSM is configured properly.

You use DSM to set up the DSM configuration group based on the partitioning of the network of systems into common file system name spaces. You use CONFIG_DSM to create or modify the membership of the DSM configuration group as in the following examples. For details on configuring DSM, see the *DSM User's Guide.*

The following example demonstrates how to expand the membership of an existing DSM configuration group to add a new system. You use this process if you have systems that have disks that are to be part of the common file system name space but the systems are not presently included in the DSM configuration group.

To create or modify the DSM configuration group, follow these steps:

1. Set up the DSM configuration group to contain those systems that will share the common file system name space. Use the MODIFY facility of the CONFIG_DSM utility to modify the membership of the DSM configuration group as in this example:

```
OK, CONFIG_DSM
[CONFIG_DSM Rev. 23.0 Copyright (c) 1989, Prime Computer, Inc.]
[Serial #IDER-QMQGA5-M821 (PRIME COMPUTER, INC.)]

Enter pathname for configuration or Quit
(default is "DSM*>CONFIG_FILES>DSM_DEFAULT.CONFIG"):


DSM*>CONFIG_FILES>DSM_RESTART.CONFIG
Configuration                : DSM_RESTART.CONFIG
Revision number              : 57
Last updated                 : 90/02/02 10:21:03
Updated by user              : SYSTEM
Updated on node              : SYSA
DSM revision number          : 1
Comment                      : PRODUCTION GROUP

--Press < RETURN > to continue:  Return

INITIAL menu.

Which of the following do you wish to perform?

  (1) MODIFY the configuration.
  (2) CHECK the configuration for consistency.
  (3) SAVE the configuration.
  (4) LIST the configuration.
  (5) HELP.

Enter option number, or press RETURN to leave this menu: 1


MODIFY the configuration.

Which of the following do you wish to modify?

  (1) Membership of the CONFIGURATION GROUP.
  (2) Membership of the ALIEN NODE GROUP.
  (3) Definition of FUNCTION GROUPS.
  (4) Definition of NODE GROUPS.
  (5) USER ACCESS DEFINITIONS.
  (6) Definition of the PRODUCT REGISTER.
  (7) LIST the configuration.
  (8) HELP.

Enter option number, or press RETURN to leave this menu: 1
```

```
MODIFY the membership of the configuration group.

The configuration group (.GROUP$) contains:
   SYSA, SYSB, SYSC, SYSD, SYSE, SYSF, SYSG, SYSH, SYSI, SYSJ

  (1) ADD node to configuration group.
  (2) REMOVE node from configuration group.
  (3) LIST the configuration.
  (4) HELP.

Enter option number, or press RETURN to leave this menu: 1

Enter node name: SYSZ
Enter node name:  Return
```

Continue pressing Return until you return to the INITIAL menu. At that point, select option (3) SAVE the configuration. From the SAVE menu, select option (2) CREATE new configuration.

```
INITIAL menu.

Which of the following do you wish to perform?

  (1) MODIFY the configuration.
  (2) CHECK the configuration for consistency.
  (3) SAVE the configuration.
  (4) LIST the configuration.
  (5) HELP.

Enter option number, or press RETURN to leave this menu: 3


SAVE the configuration.

How do you wish to save the configuration?

  (1) WRITE to temporary configuration.
  (2) CREATE new configuration.
  (3) UPDATE existing configuration.
  (4) LIST the configuration.
  (5) HELP.

Enter option number, or press RETURN to leave this menu: 3

Warning from DSM_ADMIN (DSM_admin-430):
     The configuration has not been validated. Either CHECK the
     configuration or SAVE it as a temporary configuration. Continuing
     will CHECK the configuration.
--Continue ? YES


The configuration has PASSED the consistency checks.

--Press < RETURN > to continue:  Return
Enter pathname for configuration: ADMIN>NEW.CONFIG
Enter comment or Quit (default is "DEFAULT CONFIG FILE"): New Config file
Configuration <TPUSER>ADMIN>NEW.CONFIG, revision 58 written.
--Press < RETURN > to continue:  Return
```

Continue pressing Return until you exit CONFIG_DSM.

2. Distribute the newly created DSM configuration file to all systems in the configuration group by using the DISTRIBUTE_DSM utility. When you distribute the new configuration file, it becomes the restart file (DSM_RESTART.CONFIG) in DSM*>CONFIG_FILES on all systems in the group.

Note

To distribute the configuration the first time or to a new system, you must install the configuration file in the restart file (DSM_RESTART.CONFIG) of each system or the new system by using either the COPY command or the File Transfer Service (FTS). See the *DSM User's Guide* for details.

Use DISTRIBUTE_DSM as in this example:

```
OK, DISTRIBUTE_DSM
[DISTRIBUTE_DSM Rev. 23.0 Copyright (c) 1989, Prime Computer, Inc.]

Enter pathname for configuration or Quit
(default is "DSM*>CONFIG_FILES>DSM_DEFAULT.CONFIG"):

ADMIN>NEW.CONFIG

Configuration                : DSM_SYSA.CONFIG
Revision number              : 80
Last updated                 : 90/02/16 17:47:21
Updated by user              : WSH
Updated on node              : SYSA
DSM revision number          : 15
Comment                      : NEW CONFIG FILE

--Press < RETURN > to continue:


INITIAL menu.

The following options are available:

  (1) DISTRIBUTE to a SINGLE node
  (2) DISTRIBUTE to ALL nodes in configuration group
  (3) REMOVE node from LOADED configuration group
  (4) LIST configuration group nodes in SPECIFIC configuration.
  (5) LIST configuration group nodes in LOADED configuration.
  (6) HELP.

Enter option number, or press RETURN to leave this menu: 2

The configuration being distributed is:

Configuration                : DSM_SYSA.CONFIG
Configuration UID            : RXHRDBQBBFGTS65
Revision number              : 80
Last updated                 : 90/02/16 17:47:21
Updated by user              : WSH
Updated on node              : SYSA
DSM revision number          : 15
Comment                      : NEW CONFIG FILE

Function DISTRIBUTE_DSM being invoked on node:  SYSA, SYSB, SYSC, SYSD,
SYSE, SYSF, SYSG, SYSH, SYSI, SYSJ, SYSZ
```

```
Node: SYSB
RESTART configuration successfully updated at 16 Feb 90 17:48:44 Friday.

Node: SYSC
RESTART configuration successfully updated at 16 Feb 90 17:48:52 Friday.

Error from DSM (DSM-1398):
     DSM configurations were generated from different bases.
     Unable to make or maintain connection to DISTRIBUTE_DSM on SYSD.

Node: SYSE
RESTART configuration successfully updated at 16 Feb 90 17:49:12 Friday.
          .          .          .
          .          .          .
          .          .          .
```

Because of the error DSM encountered above related to DSM configurations generated from different bases, it is necessary to use the COPY command to distribute the configuration file to SYSD this time.

```
OK, COPY ADMIN>NEW.CONFIG <SYSDO>DSM*>CONFIG_FILES>DSM_RESTART.CONFIG -NQ -RPT
"ADMIN>NEW.CONFIG" copied to "<SYSDO>DSM*>CONFIG_FILES>DSM_RESTART.CONFIG".
OK,
```

3. After distributing the new DSM configuration file by using DISTRIBUTE_DSM or COPY, restart DSM on each system in the common file system name space. To restart DSM, use the STOP_DSM command and then issue the START_DSM command at the supervisor terminal of each system.

4. Configure unsolicited message handling (UMH) on all systems in the DSM configuration group so that all Name Server error messages (product NAME_SERVER) are displayed at the supervisor terminal and go to a DSM system log such as PRIMOS.LOG. You configure UMH by using the CONFIG_UM command.

```
OK, CONFIG_UM NAME_SERVER -CREATE
[CONFIG_UM Rev. 23.0 Copyright (c) 1989, Prime Computer, Inc.]
Prime Product: NAME_SERVER
Prime Product:  Return
Customer Product:  Return
Severity: -ANY
Destination: DISPLAY -USER SYSTEM -FORMAT FULL
Destination: LOGGER DSM*>LOGS>PRIMOS>PRIMOS.LOG
Do you wish to edit this selection? NO
Configuring NAME_SERVER on SYSA
Completed OK
OK,
```

THE START_NAMESERVER COMMAND

Put the START_NAMESERVER command in the PRIMOS.COMI file of each Rev. 23.0 system in the DSM configuration group. The command must follow the START_DSM command because the Name Server must query DSM to determine the members of the DSM configuration group. START_NAMESERVER must also follow the START_NET command because the Name Server needs the network to replicate the root directory and the GMT. Starting the Name Server should also be after local disks are added but prior to any ADDISK commands for remote disks, if any are needed for ldev syntax.

Caution

If you include ADDISK commands in your PRIMOS.COMI file to start remote disks before the START_NAMESERVER command, the Name Server will not start. In addition, if you include ADDISK commands to start remote disks immediately after the START_NAMESERVER command, the Name Server may not have fully started and may abort as a result.

For example, your PRIMOS.COMI file may look like this:

```
CONFIG -DATA CONFIG
ATTACH CMDNCO                          /* Make sure we're at CMDNCO
COMO PRIMOS.COMO -NTTY                 /* Record this startup
STI -TZ 0500 -DLST YES                 /* Set Universal Time
START_DSM                              /* Start DSM early
ADD 52460 121060 3160 61560 3660 4663  /* Add local disks
START_NET                              /* Start the network
START_NAMESERVER                       /* Start the Name Server
```

Disks on other systems within the common file system name space become visible some time after the Name Server is started, depending on the network size. The Name Server on each system must broadcast its local disk information to every other system in the DSM configuration group in order to create the GMT.

If you intend to add portals by including the ADD_PORTAL command in the PRIMOS.COMI file, give the Name Server time to initialize before including ADD_PORTAL. You can do this by including some intervening command between START_NAMESERVER and ADD_PORTAL, such as a SHARE command.

For additional information, see the discussion in Chapter 5 concerning starting and stopping the Name Server process.

ADDING A NEW SYSTEM TO THE CONFIGURATION

When you want to add a new system to the common file system name space, it is only necessary to complete these items:

- Modify the network configuration to include the new system
- Modify the DSM configuration to include the new system
- Add the START_NAMESERVER command to the new system's PRIMOS.COMI file if the new system is at Rev. 23.0

Modifying the Network Configuration

Use this procedure to modify the network configuration:

1. Use CONFIG_NET to update the network configuration file.
2. Copy the network configuration file to all systems in the group, including the new system.
3. Restart the network on all systems.

Modifying the DSM Configuration

Use this procedure to modify the DSM configuration

1. Using CONFIG_DSM at one of the existing systems in the group, select (1) MODIFY the configuration from the initial menu, select option (1) Membership of the CONFIGURATION GROUP and add the new system as in the previous example.
2. Use the COPY command to copy the updated configuration file to DSM_RESTART.CONFIG in DSM*>CONFIG_FILES on the new system.
3. Use DISTRIBUTE_DSM to update the the DSM restart configuration files on all other systems in your DSM configuration group.
4. Configure UMH on the new system.
5. Stop and restart DSM on all machines in the group.

REMOVING A SYSTEM FROM THE CONFIGURATION

When you want to remove a system from the common file system name space, it is only necessary to complete these items:

- Modify the existing network configuration file to remove the system from the network
- Modify the DSM configuration file to remove the system from the DSM configuration group

Modifying the Network Configuration File

Follow these steps to modify the network configuration:

1. Use CONFIG_NET to modify the network configuration file.
2. Copy the new network configuration file to all systems in the network.
3. Stop and restart the network on all systems.

Modifying the DSM Configuration File

Follow these steps to modify the DSM configuration:

1. Using CONFIG_DSM at one of the existing systems in the group, select (1) MODIFY the configuration. from the initial menu and select option (1) Membership of the CONFIGURATION GROUP. and remove the system name.
2. Use DISTRIBUTE_DSM to update the DSM configuration files on all other systems in your DSM configuration group.
3. Stop and restart DSM on all systems to remove the system from the group.

5

OPERATION OF THE NAME SERVER

MONITORING THE COMMON FILE SYSTEM NAME SPACE

There is a new command at Rev. 23.0, LIST_MOUNTS, that allows you to list the partitions that make up the common file system name space. Two existing commands for listing partitions, LIST_DISKS and STATUS DISKS, behave differently at Rev. 23.0. You may use another existing command, LD, to list the contents of the root directory.

LIST_MOUNTS: LIST_MOUNTS allows you to

- List all disk partitions and portals in the common file system name space
- Correlate pathnames or mount-point directory names that identify file system objects to system names and disk names that identify the system and disk containing the file system objects

LIST_MOUNTS has options that allow you to select subsets of disks and portals based on specific systems or pathnames. For details, see the *Operator's Guide to System Commands.*

STATUS DISKS: STATUS DISKS lists the local disk table containing the disk partitions that are explicitly added with the ADDISK command on the local system. STATUS DISKS is *no longer* the way to determine which disks can be accessed or referenced. Use LIST_MOUNTS instead.

If the Name Server is running, remote disks that are in the local disk table are used only for converting ldevs to disk names. If a remote disk is in the local disk table as a result of an ADDISK command, this does not mean that it can be accessed or referenced. The remote disk must also be in the root directory and only the Name Server can put it there.

If a local disk is added with a mount point specified with the ADDISK -MOUNT_PATH option, the disk appears in the disk table but is *not* accessible by using a pathname beginning with the disk name. It must be accessed using the pathname of the mount point. For example, attaching to <ENGINEERING_3 from the example under LD, below, attaches you to the MFD of the partition mounted at <ENGINEERING_3.

Note

The STATUS DISKS display includes notes at the end of the display to remind you that the display may not include all disks that you have access to if you add disks with the -MOUNT_POINT option or if the Name Server is running.

LIST_DISKS: The other existing command is LIST_DISKS. LIST_DISKS also uses the local disk table, as opposed to using the GMT.

Similar to STATUS DISKS, LIST_DISKS can *not* be used to determine which disk partitions can be accessed or referenced. Use LIST_MOUNTS instead. Like STATUS DISKS, LIST_DISKS lists only the disks in the local disk table, that is, the local disks and any remote disks that were added using ADDISK.

LD: You may attach to the root directory and use the LD command to list the root directory contents. Each entry in the root directory represents a specific disk partition. For example:

```
OK, A <
OK, LD

< (LU access)

715 Directories.

ACCOUNTING        ADMINISTRATION   APPLICATIONS     ARCENA
ARCHIV            AUX12            BASEBALL         BOMDEC
CADD2A            CKLST            CMPEWN           COMMAND_DEVICE
COMDEV_SYSTEM_A   COMPRJ           CORP             COGR11
COGR20            COGR21           CPDOCUMENTATION  CQWGR1
CRONY             CSA&F            CSVINT           CSVPRJ
CSVTST            DATABASE_1       DATABASE_2       DATABASE_3
DBGRPA            DBGRPB           DBGRPC           DBGSRC
DOYSRC            DOYTST           DRAFTING_1       DRAFTING_2
DRAFTING_3        DSEED            DUMPYS           EMD123
EMD345            ENGINEERING_1    ENGINEERING_2    ENGINEERING_3
--More--Q
```

Use of the ADDISK command determines the names of the directories in the root.

- If you do not use the -MOUNT_PATH option of ADDISK, the name of the directory in the root is the name of the partition.
- If you do use the -MOUNT_PATH option, the name of the directory is the name you specify with the -MOUNT_PATH option.

With the -MOUNT_PATH option, you can specify a mount-point directory name of up to 32 characters. The pathnames of file system objects on partitions added with the -MOUNT_PATH option do not contain the name of the partition; you use the mount-point pathname instead. For example, the file ACCOUNTS in the directory USERS on the partition CPDOC that was mounted with the mount point pathname of CPDOCUMENTATION becomes, as a fully-qualified pathname, <CPDOCUMENTATION>USERS>ACCOUNTS.

USING THE NAME SERVER

This section discusses the operation of the Name Server process including starting, stopping, forcing updates of, and troubleshooting the Name Server and potential name conflicts and what to do about them.

Distributed Operation of the Server

Each Rev. 23.0 system in the common file system name space, as defined by the DSM configuration group, has a Name Server and together these Name Servers replicate the root directory on all Rev. 23.0 systems in the name space.

There is no master copy of the root directory; the various copies on the systems in the DSM configuration group are updated concurrently. The Name Servers automatically broadcast the information about the common file system name space to all other Rev. 23.0 systems in the DSM configuration group.

Name Server Assumptions About Its Environment

To summarize the recommended way to set up operation of the Name Server process to operate with a common file system name space, configure the network such that all systems in the DSM configuration group

- Are configured in PRIMENET
- Allow remote file access (RFA) among each other
- Do not use forced user validation (FUV) among each other

Configure DSM such that

- All systems that are to share in the common file system name space are in the same DSM configuration group.
- The DSM configuration group is identical on all systems in the group, that is, you centrally create the DSM configuration file and distribute it by using DISTRIBUTE_DSM.
- Unsolicited message handling (UMH) has been set up for the Name Server processes on each system.

Relationship Between the Name Server and PRIMOS Revisions

The following are the relationships between the Name Server and PRIMOS revisions.

- The set of systems in the same DSM configuration group can be at any supported PRIMOS revision.
- The Name Server can be started only on Rev. 23.0 systems.
- If the Name Server is started on one system in the DSM configuration group, it should be started on all systems in the group that are running Rev. 23.0.

Starting the Name Server

You start the Name Server by issuing the START_NAMESERVER command at the supervisor terminal, either manually or by including the command in the PRIMOS.COMI file.

If you want to start the Name Server, you should not use the ADDISK command to add any remote disks *before* you use the START_NAMESERVER command for the first time. If you do add remote disks before you start the Name Server, the Name Server will not start.

The Name Server process is started by the START_NAMESERVER command and the process takes some time to initialize. Therefore, you should not follow the START_NAMESERVER command immediately with the ADDISK -ON *nodename* command to add a remote disk. If the remote disk is added before the Name Server starts, the Name Server aborts. In general, it should be unnecessary to add remote disks with the ADDISK command when the Name Server process is running.

If you use the ADDISK command to add a remote disk *after* starting the Name Server, an entry appears in the local disk table for the remote disk. However, in order that users can access the remote disk, the disk name must be put in the root directory and in the GMT by the Name Server. The Name Server puts the disk name in the root directory and the GMT when that disk is on a system that is part of the same DSM configuration group. Thus, only disks in your DSM configuration group (that is, your name space) are accessible. To access disks in other name spaces, you must use portals.

Caution

Do not shut down remote disks and start the Name Server unless your system is currently part of a DSM configuration group running the Name Server. If your system is not part of such a group and you shut down remote disks and start the Name Server, you cannot make the remote disks available to your system again unless you stop, reconfigure, and restart DSM. In the worst case, you have to cold start if you cannot reconfigure DSM because that would result in overlapping name spaces. Adding remote disks with the ADDISK command does *not* make those disks available to your system when the Name Server is running.

Stopping the Name Server

You can stop the Name Server with the STOP_NAMESERVER command. The Name Server process is generally idle unless you shut down or add a partition to your local system, in which case your local Name Server must broadcast that information to other systems in your common file system name space. If there are pre-Rev. 23.0 systems in your name space, the Name Server polls those systems at the retry interval set by the UPDATE_NAMESERVER command.

It is not necessary to stop the Name Server. If you do stop the Name Server, any changes you make to the organization of your file system name space, such as shutting down and adding disks or adding portals, are not broadcast to the other systems in your DSM configuration group until you restart the Name Server.

Invalid Root or GMT Entries: In an extreme case, you can use the -REINIT option of the START_NAMESERVER command to remove all remote partitions and portals from the local system. Under ordinary circumstances, however, you should not use the -REINIT option. This option causes all remote partitions to disconnect, effectively shutting them down so all open file units and attach points of local users on those partitions are lost. Remote partitions are reconnected when the Name Server restarts and communicates with other Name Servers.

The primary use of the -REINIT option is to remove any information that may be present and that is no longer valid from the local copy of the root directory and the GMT.

The -REINIT option purges only the local system of remote entries. These entries reappear immediately when the local Name Server communicates with the other Name Servers in the DSM configuration group if the entries have not been purged from all systems. Use of -REINIT to purge the entries from all systems is a drastic step which should be done only if absolutely necessary and only as a last resort.

Forcing Updates From the Name Server

Name Server updates generally take less than a minute in a large network. Whenever a disk or a portal is added or removed on the local system, the local Name Server is notified and it sequences through the list of systems in the DSM configuration group and sends their Name Server the update information.

The time it takes to update all systems depends on such things as the number of systems, the speed of the network, and the speed of CPU. The time is generally in the range of a few seconds per system to be updated.

The following can interfere with the speed of updates:

- Not all systems in the DSM configuration group may have a Name Server because they may be running an earlier revision of PRIMOS.
- If there are not Name Servers on other systems, the systems with Name Servers can only determine when changes are made in these other systems by repeatedly checking, or polling, them.
- If the network is down or having intermittent failures, updates may not get through.

Because of the need to deal with polled systems and possible network failures, the Name Server retries update operations periodically if necessary. The default retry time is 10 minutes.

The UPDATE_NAMESERVER command allows you to set the retry time for the Name Server. You do this with the -RETRY option which sets the amount of time the local Name Server waits before retrying failed operations or polling pre-Rev. 23.0 systems. Using a small value for the retry time makes the Name Server very responsive to these cases but causes the Name Server to use more resources.

This command also allows manual forcing of individual updates. If you do not use any options with the UPDATE_NAMESERVER command, the local Name Server sends out the current state of its replicated root directory and GMT to all systems in the DSM configuration group without waiting for the current retry time to elapse.

Manual updates of the Name Server are useful when there are disk changes on polled (pre-Rev. 23.0) systems and the newly added disk should be available to other systems as soon as possible. Manual updates may also be useful after an error situation, such as an ISC network server malfunction that prevents the Name Server from working properly, has been corrected.

You can also use the -REMOTE option of the UPDATE_NAMESERVER command to send the replicated root directory and GMT to a particular system.

The UPDATE_NAMESERVER command and its options are described in detail in the *Operator's Guide to System Commands.*

Name Conflicts

Name conflicts occur when multiple systems independently mount disk partitions with the same name and add them to the root directory. The Name Server allows this to happen because it is more important to allow a disk to be added and possibly conflict with another remote disk than it is to disallow adding the disk until some time in the future when a network-wide check for name uniqueness can be made.

Ways That Name Conflicts Can Occur: Name conflicts can occur in the following ways:

- Multiple disks are independently added to the root with the same mount point pathname on different systems at the same time.

  If a disk with the name DISK1 is added to one system and another disk with the same name of DISK1 is added to a second system before the Name Server updates the second system, a name conflict occurs.

  If the network is not running and multiple disks are added with the same mount point pathname on different systems, a name conflict occurs when the network starts.

- A disk, such as a removable disk or a dual-ported disk, is moved from one system in the DSM configuration group to another system. If a disk is explicitly shut down, the Name Server broadcasts this information to all systems in the DSM configuration group, the entry is removed, and there is no name conflict.

  In some cases, moving a disk can cause name conflicts to occur.

If a system is shut down with the SHUTDN ALL command, the local Name Server does not broadcast this information and all this system's disks remain in the root directory of the other systems in the DSM configuration group. Name conflicts occur if these disks are then moved to other systems in the group.

If the disk came from a polled pre-Rev. 23.0 system in the same DSM configuration group and is put on a Rev. 23.0 system in the group, that disk is listed as being on the polled system until the Name Server polls again.

If a system halts, all other systems continue to assume the disks reside on the halted system, thus causing conflicts if the disks are moved.

When Name Conflicts Occur: The following things happen when name conflicts occur.

- Error messages describing the name conflict are sent by DSM. DSM UMH should be configured to display the messages at the supervisor terminal.
- The error messages are displayed on those Rev. 23.0 systems on which the conflicting disks are located.
- The systems containing the conflicting disks keep their respective conflicting disks in their name space. All other systems keep the first known disk in their name space.

The systems remain in this state until you change the situation. Error messages continue to be displayed each time the Name Server polls or retries.

What to Do about Name Conflicts: If a name conflict is due to multiple independently added disks, all but one of the conflicting disks should be removed from the common file system name space. The simplest way to remove a disk is to shut down the disk with the SHUTDN command and then re-add the disk with a different mount point pathname by using the -MOUNT_PATH option. It is not necessary to change the name of the disk in this case. Alternatively, you could shut the disk down and then re-add it using the -RENAME option of either ADDISK or SHUTDN to rename it. (See the *Operator's Guide to System Commands* for descriptions of the ADDISK and SHUTDN options.)

A more drastic way to remove the disks is to remove the systems with the conflicting disks from the current DSM configuration group.

Once all the conflicting disks have been removed, the Name Servers on all systems in the DSM configuration group automatically update the GMT so that all the systems reference the remaining disk.

If the name conflict is due to disks that have been moved from one system to another in the same DSM configuration group without first being explicitly shut down, such as a dual-ported disks or removable disks, use the ADDISK -REPLACE option to force all systems in the DSM configuration group to reference the new location for the disk.

If the name conflict is due to movement of a disk from a polled (pre-Rev. 23.0) system to a Name Server system, you can use the ADDISK -REPLACE option or you can use the UPDATE_NAMESERVER command to remove the conflict.

Effect of an Incorrect DSM Configuration

If the DSM configuration groups are not identical among systems, a local Name Server filters the name space to conform with the local system's view of the name space boundary in the following way.

- When one system receives an update from another system, if the update contains entries for disks or portals on a system that is not in the local DSM configuration group, the entries are deleted.
- Entries are unaffected for those systems that are in the local DSM configuration group but that are not in the DSM configuration group of the system sending the update.

If you stop DSM, nothing happens to the Name Server since the Name Server process reads the DSM configuration file only when the Name Server process starts or when you restart DSM. If you then restart DSM, the Name Server again filters the name space to conform with the local system's view of the name space boundary in this manner:

- The local Name Server checks to see if the list of systems in the DSM configuration group has changed since DSM was last running.
- If the list of systems is the same, the Name Server does nothing except possibly send out missed updates.
- If the list of systems is different, the Name Server changes the common file system name space to reflect the new configuration, depending on the change:
    - If a system has been removed from the list, the Name Server removes from the local system's list any disk or portal entries that originated from the removed system.
    - If a system has been added to the list, the local Name Server sends updates to the new system and the local system gets updates from the new system.
    - If a system remains in the list, disks and portals from that system are unaffected.

Filtering the Name Server updates due to mismatched DSM configuration groups consumes resources and thus should be avoided by making the DSM configuration groups identical by using DISTRIBUTE_DSM. Refer to the 4 and to the *DSM User's Guide.*

Troubleshooting the Name Server

DSM does the error reporting. The location to which the error messages go depends on how DSM unsolicited message handling (UMH) is configured on your system. Prime recommends that the messages go to the supervisor terminal and to the DSM logs. You use CONFIG_UM, as described in Chapter 4, to set the destination of messages from product NAME_SERVER. You can also use CONFIG_UM to determine these destinations.

To diagnose Name Server problems, access the DSM logs. To examine the messages in the DSM logs, use DISPLAY_LOG as in the following example.

```
OK, DISPLAY_LOG DSM*>LOGS>PRIMOS>PRIMOS.LOG -PROD NAME_SERVER
[DISPLAY_LOG Rev. 23.0 Copyright (c) 1989, Prime Computer, Inc.]

*** Message from Prime product NAME_SERVER, user NAME_SERVER on SYSA
    (Severity Warning, occurred at 13 Feb 90 16:01:12 Tuesday)
 (DSM system) SYSZ is not in the PrimeNet configuration.
                    .           .           .
                    .           .           .
                    .           .           .
```

The product NAME_SERVER messages that DSM displays are listed and explained in Appendix A.

The Name Server maintains private, detailed logs for the Name Server process in the directory SERVERS*>NAME_SERVER on each Rev. 23.0 system. In that directory, the file NAME_SERVER.COMO contains a list of the systems in the common file system name space.

In the same directory, the files NAMSVC.*date.time* log detailed information from the server about unusual events. These files have these characteristics:

- The files rotate among three files with a new log created each time the Name Server is started or the current log file becomes full.
- The events are not necessarily serious problems, but may only be informational.
- In general, the messages in these files are for a representative from your Customer Support Center.

Troubleshooting Checklist: To troubleshoot the Name Server process, you should check the following items:

- Is the network configured properly?
    - Does the network contain all the systems that are in the DSM configuration group?
    - Is RFA enabled?
    - Is there FUV between systems?
- Is the network up and running?
- Is the ISC_NETWORK_SERVER running?
- Is DSM configured properly?
    - Does the DSM configuration group contain all the systems that are to share the common file system name space?
    - Is unsolicited message handling (UMH) configured properly?
- Is DSM running before the Name Server is started?

- Is the Name Server running?
    - Was the -ON option of ADDISK used before the START_NAMESERVER command was issued, that is, are there remote disks in the disk table before the Name Server starts?
- Did you use the ARID command at the supervisor terminal before starting the Name Server? The Name Server inherits the remote ID from the supervisor terminal and when the Name Server on your system attempts to communicate with other Name Servers, the communication fails because the ID is unknown to the other Name Server. FUV can be configured between Rev. 23.0 systems running the Name Server but you should not use the ARID command for the Name Server process or at the supervisor terminal. However, it is recommended that you do not configure FUV on systems in the same DSM configuration group.

6

NETWORK EVENT MESSAGES

This chapter describes the network event messages that are produced by PRIMENET and logged by the Distributed Systems Management (DSM) facility. You can display these messages using DSM's DISPLAY_LOG command. The following is a typical entry in a log displayed by DISPLAY_LOG:

```
*** Message from product PRIMENET, generated by SYSTEM_MANAGER on SYSA
    (Severity Information, occurred at 16 Nov 89 08:31:12 Thursday)
TOKEN INSERTED INTO THE RING NETWORK
```

The first line of the entry indicates the system on which the network event was logged (SYSA, in this example). The second line indicates the severity level of the message (Information, in this example). These two lines are followed by the network event message.

Network event messages are designed primarily for support personnel and engineers to use in diagnosing persistent problems. Messages that were implemented prior to Rev. 21.0 have identifying message IDs beginning with the prefix **NETWORK_** (for example, **NETWORK_RESET**). You can use the -MSGID option of DISPLAY_LOG to display these messages. Other DISPLAY_LOG options also allow you to select and display messages of particular types, origins, and severities. For details about how to collect and display network event messages (including information about severity levels), refer to the *DSM User's Guide.*

This chapter includes three messages that are not specifically network event messages, and whose message IDs therefore do not begin with **NETWORK_**. The message IDs of these messages are **POWERF** (Power fail check), **SHUTDN** (Shutdown by an operator), and **WARM** (Warm start).

Some of the messages listed in this chapter record events such as warm starts, cold starts, and system shutdowns. Most messages record communications errors. The problems that cause network event messages can range from sporadic or one-time problems that are resolved automatically by the software, to serious problems that cause a controller to shut down or that stop the network on a system.

If a PRIMENET communications line goes down, PRIMENET stops on your system, or another network problem occurs, take the following steps:

1. Save the most recently generated network event messages (for example, all messages generated that day). The *DSM User's Guide* describes how to monitor, save, and purge log files.
2. If PRIMENET has stopped on your system, issue the START_NET command to restart it.
3. Wait to see if the problem resolves itself. For example, noise on a line may cause temporary problems that are resolved automatically and do not recur.
4. If the the problem recurs, a device shuts down, or PRIMENET stops again, issue the START_NET command to restart PRIMENET.
5. If the problem persists, contact your Customer Service Representative and have the most recent network event messages available.

The first section below lists and describes the network event messages you can display using DISPLAY_LOG. Each message is followed by its message ID. At the end of the chapter is a list of the circuit states that are displayed in some messages.

NETWORK EVENT MESSAGE DESCRIPTIONS

The following is an alphabetical list of the network event messages. Each message is followed by its message ID and a brief description.

`Bad CRC detected by MDLC` *device*`; device shut down`

Message ID: **NETWORK_PKTFLT**

An unusual software error has occurred. The device indicated by *device* has been shut down. Restart PRIMENET on the system that reported the error.

`Bad status buffer pointers detected in SLCCMP`

Message ID: **NETWORK_PKTFLT**

An unusual software error has occurred. Restart PRIMENET on the system that reported the error.

`Bad status buffer pointers detected in SLCFC`

Message ID: **NETWORK_PKTFLT**

An unusual software error has occurred. Restart PRIMENET on the system that reported the error.

Bad status detected from MDLC *device*; device shut down

Message ID: **NETWORK_PKTFLT**

An unusual error has occurred. The device indicated by *device* has been shut down. Restart PRIMENET on the system that reported the error.

Call collision in loopback mode - state 14 (X$IPKT); Network shut down

Message ID: **NETWORK_PKTFLT**

An unusual error has occurred. PRIMENET has been stopped on the system that reported the error. Restart PRIMENET on this system.

Call collision in loopback mode (X$IPKT); Network shut down

Message ID: **NETWORK_PKTFLT**

An unusual error has occurred. PRIMENET has been stopped on the system that reported the error. Restart PRIMENET on this system.

CIRCUIT RESET - LOCALLY ORIGINATED - *line*
CIRCUIT STATE: *state* CAUSE: *cause*

Message ID: **NETWORK_RESET**

A virtual circuit was reset. The reset originated on the local system. *line* indicates the link carrying the virtual circuit. *state* is the virtual circuit state. (Circuit states are described in the last section of this chapter.) *cause* is one of the following:

DTE RESET
OUT OF ORDER
REMOTE PROCEDURE ERROR
LOCAL PROCEDURE ERROR
NETWORK CONGESTION

cause may also be a word of octal data indicating the cause of the reset.

CIRCUIT RESET - LOCALLY ORIGINATED - *line*
CIRCUIT STATE: *state* CAUSE: *cause*
DIAGNOSTIC: *diag_code*

Message ID: **NETWORK_RESET**

A virtual circuit was reset. The reset originated on the local system. *line*, *state*, and *cause* are described above. *diag_code* is an octal diagnostic code.

CIRCUIT RESET - LOCALLY ORIGINATED - *line*
CIRCUIT STATE: *state* CAUSE: *cause* (HEX)
DIAGNOSTIC: *diag_code* (HEX)
FLAGS: *flags* (OCT) FLOW: *flow* (OCT)
SEQ # ACKED: *n* SEQ # OUT: *m*
SEQ # EXPECTED: *p*

Message ID: **NETWORK_RESET**

A virtual circuit was reset. The reset originated on the local system. *line* and *state* are described above. *cause* and *diag_code* are hexadecimal values indicating the cause of the reset and the diagnostic code, respectively. *flags* is an octal code used internally to control the virtual circuit. *flow* is an octal encoding of the current packet-level receive and send variables. *p* is the receive variable, which is the sequence number expected in the next incoming packet. *m* is the send variabie, which is the sequence number to be sent in the next outgoing packet. *n* is the sequence number of the last packet received and acknowledged.

CIRCUIT RESET - REMOTELY ORIGINATED - *line*
CIRCUIT STATE: *state*

Message ID: **NETWORK_RESET**

A virtual circuit was reset. The reset originated on the remote system. *line* and *state* are described in the previous entry.

COLD START - PRIMOS REV *revision*

Message ID: **NETWORK_COLD**

PRIMOS was cold started. *revision* is the PRIMOS revision number.

ICS.00 (X.25) DECONFIGURE CODE WORD NOT QUEUED FOR LOGICAL LINE *line*

ICS.01 (X.25) LOGICAL CONNECTION DELETED FOR LOGICAL LINE *line*

ICS.02 (X.25) LOGICAL CONNECTION NOT BROKEN FOR LOGICAL LINE *line*

ICS.03 (X.25) LCAD1_ NOT FOUND IN LCB FOR LOGICAL LINE *line*

ICS.04 (X.25) LOGICAL CONNECTION LOST FOR LOGICAL LINE *line*

ICS.05 (X.25) FLUSH TIMEOUT FOR LOGICAL LINE *line*

ICS.06 (X.25) ILLEGAL FLUSH COMPLETE FOR LOGICAL LINE *line*

ICS.07 (X.25) SYNCHRONOUS LINE NOT ASSIGNED FOR LOGICAL LINE *line*

ICS.08 (X.25) UNIDENTIFIABLE ERROR FOR LOGICAL LINE *line*

ICS.09 (X.25) LINE NOT DEFINED: *line*

ICS.10 (X.25) OVERSIZE FRAME RECEIVED FOR LOGICAL LINE *line*

Message IDs: **NETWORK_ICS_00** through **NETWORK_ICS_10**

The above messages indicate problems with the Intelligent Communications Subsystem, Model 1 (ICS1), Model 2 (ICS2), or Model 3 (ICS3). *line* indicates the communication link to which the message applies.

ICS.20 (X.25) INVALID COMMAND TO IBC FOR LOGICAL LINE *line*

ICS.21 (X.25) INVALID PROTOCOL ID FOR LOGICAL LINE *line*

ICS.22 (X.25) LAC BUS: UNMAPPED LINE INTERRUPT FOR LOGICAL LINE *line*

ICS.23 (X.25) LAC BUS: ADDRESS PARITY ERROR FOR LOGICAL LINE *line*

ICS.24 (X.25) LAC BUS: DATA PARITY ERROR FOR LOGICAL LINE *line*

ICS.25 (X.25) LAC BUS: PARITY ERROR ON IA CYCLE FOR LOGICAL LINE *line*

ICS.26 (X.25) UNMAPPED LINE ON DMX SCAN FOR LOGICAL LINE *line*

ICS.36 (X.25) BISYNC FRAMING ERROR FOR LOGICAL LINE *line*

ICS.37 (X.25) STATUS BUFFER OVERFLOW FOR LOGICAL LINE *line*

Message IDs: **NETWORK_ICS_20** through **NETWORK_ICS_37**

The above messages indicate problems with the Intelligent Communications Subsystem, Model 2 (ICS2) or Model 3 (ICS3). *line* indicates the communication link to which the message applies.

IPQNM problem detected in SLXDIM

Message ID: **NETWORK_PKTFLT**

An unusual error has occurred. Restart PRIMENET on the system that reported the error.

LEVEL III PROTOCOL DOWN - *line*

Message ID: **NETWORK_HOSTDN**

The X.25 level 3 protocol is down on the link indicated by *line*.

Level 3 network received a diagnostic packet

Message ID: **NETWORK_DIAPKT**

The second line of this message is one of the following, indicating the source of the diagnostic packet -- a system acting as Data Terminal Equipment (DTE), or a PSDN with Data Network Identification Code (DNIC) *dnic*:

```
This packet was sent by a DTE
DNIC of the PDN is: dnic
```

The third and fourth lines indicate the link over which the diagnostic packet was received:

```
The controller number is: controller (OCT)
The line number is: line (OCT)
```

The fifth line contains one of the following X.25-related messages:

```
No additional info, dcode is: code (OCT)
Packet not allowed, dcode is: code (OCT)
Packet on an unassigned Lchannel: code (OCT)
Packet too short, dcode is: code (OCT)
Invalid GFI, dcode is: code (OCT)
Timer expired, dcode is: code (OCT)
Timer expired for CLEAR INDICATION: code (OCT)
Timer expired for RESET INDICATION: code (OCT)
Timer expired for RESTART INDICATION: code (OCT)
The diagnostic code is: code (OCT)
```

The sixth line contains the words `Diagnostic explanation` followed by an octal string.

LOCAL PROCEDURAL ERROR CAUSING CLEAR - *line*

Message ID: **NETWORK_LPE**

A local procedure error caused the clearing of a circuit on the link indicated by *line*.

LOGICAL CONNECTION FAILURE TO LHC controller dequeue error: *err*

Message ID: **NETWORK_LCFAIL**

The PRIMENET server could not connect with the software on the LAN Host Controller (LHC) board. *err* is the controller error code.

LOGICAL CONNECTION FAILURE TO LHC controller enqueue error: *err*

Message ID: **NETWORK_LCFAIL**

The PRIMENET server could not connect with the software on the LAN Host Controller (LHC) board. *err* is the controller error code.

MDLC *device* missed OTA and overflowed tumble tables; device shut down

Message ID: **NETWORK_PKTFLT**

An unusual error has occurred. The device indicated by *device* has been shut down. Restart PRIMENET on the system that reported the error.

MDLC *device* missed OTA in call to SLIOC from SLCCMP; device shut down

Message ID: **NETWORK_PKTFLT**

An unusual error has occurred. The device indicated by *device* has been shut down. Restart PRIMENET on the system that reported the error.

MDLC *device* missed OTA in call to SLIOC from SLCNET; device shut down

Message ID: **NETWORK_PKTFLT**

An unusual error has occurred. The device indicated by *device* has been shut down. Restart PRIMENET on the system that reported the error.

MDLC *device* missed OTA in call to SLIOC from SLCNSB; device shut down

Message ID: **NETWORK_PKTFLT**

An unusual error has occurred. The device indicated by *device* has been shut down. Restart PRIMENET on the system that reported the error.

NETDMP CALLED AT *addr1 addr2* (OCT)
DATA = *data1 data2 data3* (OCT)

Message ID: **NETWORK_NETDMP**

A network software problem has occurred at the address indicated by the octal numbers *addr1* and *addr2*. The routine NETDMP was called and was asked to dump the three octal DATA words (*data1, data2, data3*).

No blocks available to send clear packet (X$CLOS); Network shut down

Message ID: **NETWORK_PKTFLT**

An unusual error has occurred. PRIMENET has been stopped on the system that reported the error. Restart PRIMENET on that system.

No room at top of transmit queue (X$GETU); Network shut down

Message ID: **NETWORK_PKTFLT**

An unusual error has occurred. PRIMENET has been stopped on the system that reported the error. Restart PRIMENET on that system.

No room on receive queue (X$RCV); Network shut down

Message ID: **NETWORK_PKTFLT**

An unusual error has occurred. PRIMENET has been stopped on the system that reported the error. Restart PRIMENET on that system.

No room on transmit queue (X$TRAN); Network shut down

Message ID: **NETWORK_PKTFLT**

An unusual software malfunction has occurred. PRIMENET has been stopped on the system that reported the error. Restart PRIMENET on that system.

NPX>R$CALL>R$CONN UNKNOWN CIRCUIT STATUS - NODE: *node* (OCT)
VIRTUAL CIRCUIT STATE (1): *state1* (OCT)
VIRTUAL CIRCUIT STATE (2): *state2* (OCT)

Message ID: **NETWORK_NPXCON**

PRIMENET has returned an unexpected status (error) code to NPX. The cause may be a hardware failure. *node*, *state1*, and *state2* are octal numbers indicating the communication link and the virtual circuit states. (Virtual circuit states are listed in the last section of this chapter.)

NPX>R$RLS ERROR IN VIRTUAL CIRCUIT CLEARING - NODE: *node* (OCT)
VIRTUAL CIRCUIT STATE (1): *state1* (OCT)
VIRTUAL CIRCUIT STATE (2): *state2* (OCT)

Message ID: **NETWORK_NPXRLS**

A problem occurred during virtual circuit clearing. The returned virtual circuit status word 2 is not one of the existing status codes. *node*, *state1*, and *state2* are octal numbers indicating the communication link and the virtual circuit states. (Virtual circuit states are listed in the last section of this chapter.)

NPX>SLAVE TRNRCV FAILED, RETURN CODE: *code* (OCT)

Message ID: **NETWORK_NPXETR**

An NPX slave TRNRCV failure was detected. *code* is the octal return code.

NPX>SLAVER UNABLE TO ASSIGN PORT, RETURN CODE: *code* (OCT)

Message ID: **NETWORK_NPXPUA**

NPX was unable to assign a port. *code* is the octal return code.

NPX>TRNRCV MASTER'S CIRCUIT WAS CLEARED - NODE: *node* (OCT)
VIRTUAL CIRCUIT STATE (1): *state1* (OCT)
VIRTUAL CIRCUIT STATE (2): *state2* (OCT)

Message ID: **NETWORK_NPXCLR**

The connection between the master and the slave has been unexpectedly broken. *node*, *state1*, and *state2* are octal numbers indicating the communication link and the virtual circuit states. (Virtual circuit states are listed in the last section of this chapter.)

NPX>TRNCV MESSAGE OUT OF SEQUENCE IN BOUNCE DETECT -
NODE: *node* (OCT), MESSAGE *sequence* (OCT), NS: *ns* (OCT)

Message ID: **NETWORK_NPXSEQ**

NPX break-detection and correction logic found a message out of sequence. NPX has failed or data have been lost in the network. *node* is an octal number indicating the communication link. *sequence* and *ns* are the expected and received sequence numbers, respectively.

NPX>TRNRCV THROTTLED ON XMIT OR RCV -
NODE: *node* (OCT), MASTER/SLAVE FLAG: *flag* (OCT)

Message ID: **NETWORK_NPXTHR**

Network buffers are too full to send or receive an NPX message. *node* is an octal number indicating the communication link. *flag* is either 0 (if the message was generated by the slave system) or 1 (if the message was generated by the master system).

NPX>TRNRCV UNKNOWN RECEIVE STATUS - NODE: *node* (OCT).
MASTER/SLAVE FLAG: *flag* (OCT). RECEIVE STATE: *state* (OCT)

Message ID: **NETWORK_NPXRCV**

PRIMENET has returned an unrecognized status (error) code to NPX. *node* is an octal number indicating the communication link. *flag* is either 0 (if the message was generated by the slave system) or 1 (if the message was generated by the master system). *state* is the unrecognized status code.

OUT OF RECEIVE BLOCKS FOR LHC controller posting error: *err*

Message ID: **NETWORK_NORCVB**

An internal buffer pool has been exhausted. *err* is the controller error code. Restart PRIMENET on the system that reported the error.

PACKET OUT OF SEQUENCE - *line* CIRCUIT STATE: *state*
SEQ # EXPECTED: *expected*, SEQ # FOUND: *found*

Message ID: **NETWORK_BADSEQ**

The system received a packet with sequence number *found*, but was expecting sequence number *expected*. *line* indicates the link on which the packet was received. *state* is the virtual circuit state. (Virtual circuit states are listed in the last section of this chapter.)

Parity error detected in SLCNET; device shut down

Message ID: **NETWORK_PKTFLT**

An unusual error has occurred. Restart PRIMENET on the system that reported the error.

Per-line sync process semaphore has wrapped in SLCCMP

Message ID: **NETWORK_PKTFLT**

An unusual error has occurred. Restart PRIMENET on the system that reported the error.

PNC HARDWARE FAILURE

Message ID: **NETWORK_RNGHRD**

A PRIMENET Node Controller (PNC) hardware malfunction has occurred. The PNC has been shut down, and the system has removed itself from the ring. Hardware diagnostic tests should be run on the PNC. One of the following additional messages is displayed:

```
DMA FAILURE
NO SKIP ON INA
NO SKIP ON RECEIVE OTA
NO SKIP ON TRANSMIT OTA
```

POWER FAIL CHECK

Message ID: **POWERF**

A power failure check has occurred.

PRIMENET BUFFER OVERFLOW -- *n* MESSAGES LOST, *m* BYTES LOST

Message ID: **NETWORK_OVERFL**

Network event messages have been lost because of a buffer overflow.

Queueing problem encountered (X$GIVU); Network shut down

Message ID: **NETWORK_PKTFLT**

An unusual error has occurred. PRIMENET has been stopped on the system that reported the error. Restart PRIMENET on that system.

RECEIVED A BAD RESTART PACKET (HDX) from SYNC line: *line* (OCT)

Message ID: **NETWORK_HDXBRS**

The system reporting the error received an invalid half-duplex restart packet over the link indicated by *line*. (An HDX restart packet is used to initiate HDX communications.) This message occurs in any of the following cases:

- The restart packet is too short (fewer than 5 bytes).
- The restart packet was corrupted.
- The calling or called system's name length is invalid.
- The incoming call is for a different system.
- The password is too short (fewer than 2 bytes).
- The incoming password does not match the expected password.
- The link was already established. The calling system went down, came up again, and initiated another call over the same line, but intended the new call for a different system.
- The calling system is not configured as an HDX site.
- The software is unable to get a transmit packet from the free pool of packets.

The calling system should retry the call. If the problem recurs, restart PRIMENET on both systems.

RING DIM OUT OF RECEIVE BLOCKS

Message ID: **NETWORK_RING2**

The available RINGNET receive buffers are temporarily exhausted because of heavy incoming traffic. Remote systems transmitting to the system reporting the error will experience delays.

RING_NODE: *node* NOT ACCEPTING XMITS.

Message ID: **NETWORK_RING3**

Remote ring node *node* is not accepting transmit packets from the local system. The local system therefore marks *node* as DOWN in the STATUS NETWORKS display. This message is followed by one of four additional informational messages:

PACKET LOST, RING DOWN

The whole ring is down.

PACKET WACKED

The specified remote node is in the ring, but is unable to receive packets from the rest of the ring. The remote node may have halted, or its incoming traffic may be extremely congested.

NODE NOT IN RING

The specified node has left the ring (the network has been shut down on it).

XMIT STAT IS: *code* (OCT)

The problem may be caused by a NAK or CRC failure (due, for example, to noise on the line or excessively long cables). *code* is the octal value of the transmit status word from the PNC.

Ring Node_ *node_number* (OCT) Receive TIMEOUT - node down

Message ID: **NETWORK_RNGTMT**

The local system has not received a packet from ring node *node_number* for three minutes. Ring node *node_number* is marked as DOWN in the STATUS NETWORK display.

RING QUEUE OVERFLOW: DIM TO LEVEL II - RECEIVE PACKET LOST

RING QUEUE OVERFLOW: LEVEL II TO DIM - TIMER PACKET LOST

RING QUEUE OVERFLOW: LEVEL II TO DIM - TRANSMIT PACKET LOST

Message ID: **NETWORK_RNGRES**

The above three messages indicate that one of the queues used to move packets to and from the Prime Node Controller Device Interface Module (PNCDIM) has overflowed. Since the queues are designed to be large enough to handle peak traffic numbers, this message indicates a software malfunction. The packet being queued is returned to the free pool and ignored.

SHUTDOWN BY AN OPERATOR

Message ID: **SHUTDN**

The operator issued a SHUTDN ALL command, causing an automatic dump of the network event log buffer.

SPURIOUS RECEIVE INTERRUPT ON PNC

Message ID: **NETWORK_RNGRCV**

A receive interrupt was issued by the Prime Node Controller (PNC) when no receive was pending. This error indicates a hardware malfunction which disconnects the PNC from the ring. It is recommended that you run hardware diagnostic tests on the PNC.

SYNC - NO STX PRECEDING ETX.
PHYSICAL LINE NUMBER IS *number* (OCT)
DEVICE ADDRESS IS *address*

Message ID: **NETWORK_SYNC2**

Packets sent over synchronous lines using Bisync (BSC) framing must begin with DLE/STX and end with DLE/ETX. A packet with an invalid format was detected on the link indicated by *number* and *address*.

SYNC RESET FOR LOGICAL LINE *line* - *code*

Message ID: **NETWORK_SYNC4**

Synchronous line *line* has been reset. *code* provides additional information, as follows:

Code	*Meaning*
1	Invalid Address: The system reporting the error received an invalid level 2 address.
2	Command Reject: The system reporting the error received an invalid command byte.
3	Invalid NR: The system reporting the error received a receive number that was out of the legal range (0 through 7) or out of sequence.
4	Invalid Response: The system reporting the error received an invalid response from the remote system.
5	Invalid NR on Reject: The remote system asked the system reporting the error to resend data that had not yet been sent once.
6	Max Number of Retries Exceeded: The system reporting the error retried a transmit the maximum number of times allowed.

If *code* is none of the above, the message UNDEFINED CAUSE is displayed.

SYNC STATUS ERROR. STATUS WORD IS *status* (OCT)
PHYSICAL LINE # IS *line* DEVICE ADDRESS IS *addr* (OCT)
NUMBER OF OCCURRENCES IS *number*

Message ID: **NETWORK_SYNC1**

An invalid status, *status*, has been reported on line *line* by the synchronous controller indicated by device address *addr*. The number of occurrences, *number*, is printed only in the case of parity errors. This problem is resolved automatically by the software.

SYNC5 - CMDR SENT FOR LOGICAL LINE *line*

Message ID: **NETWORK_SYNC5**

An X.25 command reject (CMDR) was sent on line *line*. (The three-octet data field showing the cause is also displayed.) This problem is resolved automatically by the software.

SYNC6 - INTERNAL LEVEL 2 ERROR FOR LOGICAL LINE *line*
ERROR CODE = *code* LINE CONTROL BLOCK ADR = *addr1 addr2* (OCT)

Message ID: **NETWORK_SYNC6**

An internal error occurred at level 2. Restart PRIMENET on the system reporting the error.

SYSTEM BLOCKS UNAVAILABLE FOR SYNC PROTOCOL MESSAGE
MESSAGE IS *message* (OCT), LOGICAL LINE NUMBER IS *line*

Message ID: **NETWORK_SYNC3**

The level 2 synchronous protocol had no buffers in which to send the type of protocol-generated message indicated by *message*. This message is displayed if an attempt is made to send a level 2 frame when the system free pool has been exhausted. Restart PRIMENET on the system reporting the error.

TOKEN INSERTED INTO THE RING NETWORK

Message ID: **NETWORK_RING1**

The software controlling the Prime Node Controller (PNC) has determined that the ring network control token does not exist, and has therefore issued a new one. This event occurs, for example, when the ring starts up and when the token is lost.

Transmit queue full (X$XMIT); Network shut down

Message ID: **NETWORK_PKTFLT**

An unusual error has occurred. PRIMENET has been stopped on the system that reported the error. Restart PRIMENET on that system.

Virtual circuit in undefined state (X$IPKT); Network shut down

Message ID: **NETWORK_PKTFLT**

An unusual error has occurred. PRIMENET has been stopped on the system that reported the error. Restart PRIMENET on that system.

`Virtual circuit in unknown state (X$CLOK); Network shut down`

Message ID: **NETWORK_PKTFLT**

An unusual error has occurred. PRIMENET has been stopped on the system that reported the error. Restart PRIMENET on that system.

`Virtual circuit not in defined state (X$CLOK); Network shut down`

Message ID: **NETWORK_PKTFLT**

An unusual error has occurred. PRIMENET has been stopped on the system that reported the error. Restart PRIMENET on that system.

`WARM START`

Message ID: **WARM**

PRIMOS was warm started.

CIRCUIT STATES

Some network event messages include a `circuit state` code. Circuit states indicate the status of the virtual circuit involved in the network event. For general information about virtual circuits and the virtual circuit status array, refer to the *Programmer's Guide to Prime Networks*. The following are the circuit state codes and their meanings:

State	*Meaning*
1	Remote login from the local system
2	Not used
3	Not used
4	Circuit being transferred
5	User data transfer
6	Outgoing local user call request pending
7	Incoming remote user call request pending
8	Outgoing local user clear request pending

9 Not used

10 Not used

11 Not used

12 Clear desired but no memory available. PRIMENET will automatically retry the clear request.

13 Not used

14 Remote logthrough to another system

15 Clear indication has been received; awaiting clear confirm

16 Call request awaiting restart

7

MISCELLANEOUS PRIMENET INFORMATION

This chapter contains PRIMENET update information on the following topics:

- START_NM and STOP_NM commands
- START_NET command
- Gateway configuration guidelines
- ACL group restriction for Remote File Access
- XLCONN and XLASGN subroutines

START_NM AND STOP_NM

Rev. 22.1 introduced two new commands, START_NM and STOP_NM, for starting and stopping the Network Management Services process, NM_SERVER. NM_SERVER manages LAN300 Host Controller (LHC300) and Intelligent Communications Subsystem, Model 3 (ICS3) controllers, providing continuous, automatic monitoring and recovery from failures. NM_SERVER also supports network services that use LHC controllers.

The START_NM command starts up NM_SERVER. It is recommended that you place START_NM toward the beginning of the PRIMOS.COMI file, after the START_DSM command but before the COMM_CONTROLLER command, as follows:

```
START_DSM
START_NM
COMM_CONTROLLER
```

This ordering has the following benefits:

- Messages are logged to the DSM logs during the COMM_CONTROLLER command, creating a permanent record of any problems.
- Polling and recovery are provided as soon as possible for any ICS3s that are already

running. The ICS3s are downline loaded during PRIMOS coldstart. They may already be running asynchronous protocols even if network products such as PRIMENET and Network Terminal Service (NTS) are not yet running.

If you do not include START_NM in the PRIMOS.COMI file, the START_NET command starts up NM_SERVER automatically when you start PRIMENET.

To display START_NM's syntax, enter

```
START_NM -HELP
```

STOP_NM logs NM_SERVER out. Services that use the LAN300 while NM_SERVER is logged out are not provided with Network Management functions. Therefore, you should issue the STOP_NM command only after the Network Management functions and the services that depend on them are stopped. For example, a procedure for stopping Network Management might be the following:

```
STOP_NET (if PRIMENET is running)
STOP_NTS (if NTS is running)
STOP_NM
```

To display STOP_NM's syntax, enter

```
STOP_NM -HELP
```

START_NET COMMAND

The START_NET command has two new command line options at Rev. 23.0:

-MAX_IND_PATHS_PERNODE ***max_paths***
-MPA ***max_paths***

-MAX_PSDNS_PERPATH ***max_psdns***
-MPS ***max_psdns***

These options allow you to limit the number and kinds of indirect network paths that START_NET stores when it starts PRIMENET.

When START_NET starts PRIMENET, it reads the network configuration file created by CONFIG_NET. START_NET stores information about the local system's direct connections to remote nodes. It also stores information about paths through the network to nodes that are not directly connected to the local system. These paths are called **multi-hop** or **indirect** paths. PRIMENET uses this stored information when making calls to remote nodes.

The MAX_IND_PATHS_PERNODE option sets *max_paths* as the maximum number of indirect paths that START_NET can store for any remote node. *max_paths* must be a

number from 0 through 4. If you do not include this option, START_NET stores a maximum of four indirect paths for each remote node. (Prior to Rev. 23.0, four was the hard coded maximum.)

The MAX_PSDNS_PERPATH option sets *max_psdns* as the maximum number of PSDNs allowed in a path to a remote node. *max_psdns* must be a number from 0 through 4. If you do not include this option, START_NET imposes a limit of four PSDNs per path.

GATEWAY CONFIGURATION GUIDELINES

The following are clarifications of two guidelines presented on page 1-13 of the Rev. 22.0 *PRIMENET Planning and Configuration Guide.*

- The book currently states, "Configure only one gateway node between a node and a PSDN connection." This guideline should read as follows: "Configure at most one gateway node between source and destination (calling and called) nodes in any path containing a PSDN connection." Thus, the path

  SYSA -- SYSB -- PSDN -- SYSC

  (containing gateway SYSB) supports communication between SYSA and SYSC; but the paths

  SYSA -- SYSB -- SYSC -- PSDN -- SYSD
  SYSA -- SYSB -- PSDN -- SYSC -- SYSD

  (containing gateways SYSB and SYSC) do not support communication between SYSA and SYSD.

- The book currently states, "If a route-through path includes a PSDN, the PSDN must support subaddressing or multiple addresses per PSDN line." This guideline should read as follows: "If a route-through path includes a PSDN, the PSDN must support either full addressing with subaddresses, or multiple addresses per PSDN line."

ACL GROUP RESTRICTION FOR REMOTE FILE ACCESS

The following note is for users, System Administrators, and Network Administrators on systems where Remote File Access (RFA) requests are generated. This note should be added to Chapter 2 of the *User's Guide to Prime Network Services* and to Chapter 3 of the *PRIMENET Planning and Configuration Guide.*

Note

An RFA request in which the remote system does not force user validation (that is, does not require a remote ID) fails if the requesting user is assigned to more than thirty Access Control List (ACL) groups.

XLCONN AND XLASGN SUBROUTINES

The following corrections apply to Chapter 4 of the Rev. 21.0 *Programmer's Guide to Prime Networks*:

- The correct spelling of the long form subroutine for establishing a virtual circuit is XLCONN. (The subroutine is incorrectly labelled XLCON in Chapter 4.)
- The following is the correct call syntax for the XLASGN subroutine:

```
DCL XLASGN ENTRY(BIT(16), CHAR(15)VAR, CHAR(15)VAR, CHAR(4)VAR,
                 CHAR(128)VAR, FIXED BIN(15), CHAR(41)VAR,
                 FIXED BIN(15), FIXED BIN(15), FIXED BIN(15),
                 FIXED BIN(15), FIXED BIN(15));

CALL XLASGN (key, tadr, tsadr, prid, udata, port, txadr, rsrvd1,
             rsrvd2, rsrvd3, count, status);
```

The *count* argument is omitted from the call syntax in Chapter 4.

APPENDICES

A

NAME SERVER MESSAGES REPORTED BY DSM

DSM logs the following messages and displays them at the supervisor terminal. The messages are listed alphabetically, sorted by the second word if the first is an article such as The. A brief explanation of each message follows the message.

The configuration failed for the (DSM ConfigGroup) NameSpace Boundaries.
The Name Server could not get the list of machines from DSM that are used to define the common file system name space. Check to see if DSM is running. If DSM is not running, restart it and then issue the START_NAMESERVER command again.

(DSM system) *system_name* is not in the PrimeNet configuration.
DSM returned a node name of a system that is not known to PRIMENET. Either add the node to PRIMENET or remove it from the DSM configuration group.

The NAME SERVER cannot be started with remote-addisks (ADDISK -ON).
You used the START_NAMESERVER command on a system that already has remote disks added. Shut down the remote disks and use the START_NAMESERVER command again to start the Name Server.

The NAME SERVER has been started.
You used the START_NAMESERVER command and the Name Server process is now running on your system.

The NAME SERVER has terminated.
Either you used the STOP_NAMESERVER command or the Name Server terminated abnormally. If you did not explicitly stop the Name Server, check the logs in SERVERS*>NAME_SERVER for information as to why it stopped.

Name-Conflict: two locations found. Duplicate pathname also exists on *system_name*.
The local system sent its version of the GMT to another system, *system_name*. In the process of reconciling the version of the GMT on the other system with the one sent by the local system, a name conflict was found. (A name conflict occurs when two systems have a disk mounted at the same place in the file system.) The name conflict is a result of a disk added to the local system. Consult the other system to determine

the conflicting mount-point pathname and remove the source of the conflict from either the local system or the other system.

Name-Conflict: two locations found.
Duplicate pathname *pathname* also exists on *system_name*.

The local system received a version of the GMT from another system, *system_name*. In the process of reconciling the local version of the GMT with the one on the other system, the local system has found a name conflict at the mount point *pathname*. (A name conflict occurs when two systems have a disk mounted at the same place in the file system.) The conflict is due to a locally added disk that conflicts with a disk added on *system_name*. The conflict must be resolved either on the local system or on *system_name*.

Possible Name-Conflict: new location found. Pathname
pathname now found on *system1*, was also on *system2*.

The local system received a version of the GMT from another system, *system_name1*. In the process of reconciling the local version of the GMT with the one on the other system, the local system has found a name conflict at the mount point *pathname*. (A name conflict occurs when two systems have a disk mounted at the same place in the file system.) The GMT of the local system indicates that mount point to be on *system2* but the GMT from the other system indicates it to be on *system1*. Check *system1* and *system2* to see if there really is a name conflict and correct it.

Possible Name-Conflict: new location found.
system_name has a duplicate pathname with conflicting location.

The local system sent its version of the GMT to another system, *system_name*. In the process of reconciling the version of the GMT on the other system with the one sent by the local system, a name conflict was found. (A name conflict occurs when two systems have a disk mounted at the same place in the file system.) The name conflict is *not* the result of adding a disk to the local system. Consult the other system to determine the conflicting mount point pathname and the system where the conflict is.

This Primos revision does not support NAME SERVER.

You used the START_NAMESERVER command on a system that has been reverted from Rev. 23.0 to pre-Rev. 23.0.

INDEX

INDEX

E

F

G

L

M

N

P

R

S

U

X

SURVEYS

READER RESPONSE FORM

Rev. 23.0 Prime Networks Release Notes
RLN10252-1LA

Your feedback will help us continue to improve the quality, accuracy, and organization of our publications.

1. How do you rate this document for overall usefulness?

   ☐ *excellent* ☐ *very good* ☐ *good* ☐ *fair* ☐ *poor*

2. What features of this manual did you find most useful?

   _____
   _____
   _____
   _____
   _____
   _____

3. What faults or errors in this manual gave you problems?

   _____
   _____
   _____
   _____
   _____
   _____

4. How does this manual compare to equivalent manuals produced by other computer companies?

   ☐ *Much better* ☐ *Slightly better* ☐ *About the same*
   ☐ *Much worse* ☐ *Slightly worse* ☐ *Can't judge*

5. Which other companies' manuals have you read?

   _____
   _____

Name:_____ Position:_____

Company:_____

Address:_____

_____

_____ Postal Code:_____

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

First Class Permit #531 Natick, Massachusetts 01760

BUSINESS REPLY MAIL

Postage will be paid by:

Prime™

Attention: Technical Publications
Bldg 10
Prime Park, Natick, Ma. 01760